MURAL Software Upgrade Guide

Version 4.0

Published: 2017-05-04

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706 USA

http://www.cisco.com

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB’s public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS” WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

MURAL Software Upgrade Guide

TOC

Upgrade Overview

MURAL provides Web-based reporting and analytics for deep packet inspection (DPI) data emerging from your network.

This document describes how to upgrade the MURAL application from release 3.9 to release 4.0.

Before You Begin

This document assumes that you have a working knowledge of Linux operating system. Before you begin the upgrade, we recommend that you review the Release Notes for 4.0.

Upgrade Package

The MURAL software upgrade package contains the following components:

- An ISO image file. For the image name and associated MD5 checksum, refer to the Release Notes.
- The XML file which is used by the master GMS node. This configuration file provides configuration settings for your setup based on the network topology.
- Any software patches that apply to the release. A complete list appears in the Release Notes.
- Management information bases (MIBs).

Pre-Upgrade Tasks

Before you begin the upgrade, perform the following tasks:

- "Taking Backups" below
- "Stopping Processes" on page 9

Taking Backups

It is recommended to take all backups before proceeding with the upgrade.

Create a folder, where all the backed-up files will be saved, on all the nodes:

```
> en
# _shell
# mkdir /data/Backup_files
```

Take a Backup of all IBs

Perform the following steps on the master name node:

1. Create a folder where all the backed up files will be saved:

   ```
   > en
   # _shell
   # mkdir /data/Backup_files
   ```

2. Take a backup of all the IBs:

   ```
   # cp -r /data/ib/inbox /data/Backup_files/
   # cp -r /data/ib/work /data/Backup_files/
   ```

3. Execute the following command:

   ```
   # hdfs dfs -get /IB /data/Backup_files/
   ```

Take a Backup of PGSQL Database

1. Execute the following commands on the master GMS node:

```
>en
# config terminal
(config)# pgsql file-path /data/Backup_files/pgsqlDB3.9_
gms.txt dbname gms_configuration backup
```

2. Execute the following commands on the master iNSTA node:

```
>en
# config terminal
(config)# pgsql file-path /data/Backup_files/pgsqlDB3.9_
rubix.txt dbname rubixdb backup
```

3. Execute the following commands on the master Name node:

```
>en
# config terminal
(config)# pgsql file-path /data/Backup_files/pgsqlDB3.9_
hive.txt dbname hivedb backup
```

Take a Backup of Active Configurations

- Execute the following commands from the `shell` prompt on all the nodes in your setup to take a backup of configurations:

```
> en
# _shell
# cli -t "en" "conf t" "show running full" > /data/Backup_
files/`hostname`_configBackup.txt
```

Here, `hostname` is the hostname of the server. For example, if the hostname of the server is NN, then the file is created as `NN_configBackup.txt`.

Take a Backup of Gateway Configurations

- Execute the following command from all the collector nodes to take a backup of gateway configurations for bulkstats:

```
# grep "bulkStatsFile" /data/Backup_files/`hostname`_
configBackup.txt > /data/Backup_files/coll_bsGW_config.txt
# echo "write memory" >> /data/Backup_files/coll_bsGW_
config.txt
```

Take a Backup of Manually Modified Files

Perform the following procedure:

1. Execute the following commands to create folders on the nodes where you want to keep backup files:

```
> en
# _shell
# mkdir -p /data/Backup_files/Patch_backup
```

2. Execute the following command on master Name node and master DPI node:

```
# cp /data/CoreJob/config/solutionConfig.json  /data/Backup_
files/
# cp /opt/catalogue/atlas/dynamicWhiteListConfig.json
/data/Backup_files/image_dynamicWhiteListConfig.json
# cp /data/ib/inbox/dynamicWhiteListConfig.json /data/Backup_
files/
# cp -r /data/patchspace/patch/3.9.rc1 /data/Backup_
files/Patch_backup/
# cp /opt/samples/yarn_conf/capacity-scheduler.xml.template
/data/Backup_files/
```

Execute the following command on the master DPI node:

```
# cp /opt/tms/apache-tomcat/apache-tomcat-
7.0.27/conf/server.xml /data/Backup_files/
```

Stopping Processes

You must stop processes as explained in the following sections before proceeding with the upgrade.

Stop Oozie Jobs

- Log into master Name node and execute the following commands:

```
> en
# _shell
# pmx subshell oozie stop jobname all
```

Note: Wait until all the jobs are stopped. To verify if all the jobs are stopped, execute the following command and ensure that no job is running:

```
# pmx subshell oozie show coordinator RUNNING jobs
```

For example:

```
[admin@MUR-COLL-2 ~]# pmx subshell oozie show coordinator
RUNNING jobs
No Jobs match your criteria!

[admin@MUR-COLL-2 ~]#
```

Stop Rubix Processes

- To disable all applications, execute the following commands on the corresponding UI nodes:

```
> en
# conf t
(config) # rubix modify-app <app-name> modify-instance 1
disable
(config) # rubix modify-app <app-name> disable
(config) # no pm process rubix launch auto
(config) # pm process rubix terminate
(config) # write memory
```

Where, <app-name> is:

- atlas
- bulkstats
- reportAtlas

- rge
- ruleEngine

Run the following command on the nodes where launcher is configured:

```
> en
# conf t
(config)# rubix delete-app launcher
(config)# write memory
```

Stop Insta Processes

- Execute the following command on both the Insta nodes:

```
> en
# conf t
(config)# pm process insta terminate
```

Stop GMS Processes

- Execute the following commands on both the GMS nodes:

```
> en
# conf t
(config)# no pm process gms_server launch auto
(config)# no pm process pgsqld launch auto
(config)# wr mem
(config)# pm process gms_server terminate
(config)# pm process pgsqld terminate
```

Stop the Collector

- Execute the following commands on all the collector nodes:

```
> en
# conf t
(config)# no pm process collector launch auto
(config)# pm process collector terminate
(config)# write memory
```

Upgrade Tasks

Perform the following steps to upgrade MURAL from release 3.9 to release 4.0:

- "Upgrade Image on the Nodes" below
- "Restart Processes on GMS Nodes" on page 14
- "Upgrade MURAL and Bulkstats Databases on Insta Cluster" on page 14

Upgrade Image on the Nodes

Upgrade the image on the following nodes:

- Standby Collector Node
- Standby NameNode
- Master and Standby node of the additional Collector
- Master and standby GMS nodes
- All the Rubix nodes including the RGE node
- Master and standby Insta nodes
- Master and standby Compute nodes

  Note: Ensure to upgrade image on Compute nodes rack wise. After upgrading the compute nodes on the first rack, verify that the nodes have come up with new image and joined the YARN cluster. Only after this verification, proceed with the upgrade on the next rack.

Perform the following steps on all the nodes:

1. Log into the master GMS node and execute the following commands:

```
> en
# _shell
# /opt/tms/bin/cli -t en 'conf t' 'username imguser password img!1234'
# /opt/tms/bin/cli -t en 'conf t' 'write memory'
```

2. Create a directory:

```
# mkdir -p /data/Mural4.0/
```

3. Copy the image at `/data/MURAL4.0/` on the Master GMS node.

4. Log into each node in your setup and execute the following commands:

```
# cli -m config
(config)# image fetch scp://imguser:img!1234@gms-
2:/data/MURAL4.0/image-atlas4.0.rc2.img
(config)# image install image-atlas4.0.rc2.img
(config)# image boot next
(config)# write memory
(config)# reload
```

When you upgrade the image on a node, wait until the image is successfully upgraded on that node before proceeding to the other node. For example, wait until the standby GMS nodes come up with the new image before running these commands on the master GMS node.

Run the following commands on all the nodes:

```
> en
> _shell
# mount -o remount,rw /
# sed -i "s/4.0.rc1/4.0.rc2/g" /etc/solution_version
# /bin/rm -f /opt/tms/java/GPEHParser.jar
```

To verify if the new image is available on a node, execute the following commands:

```
host [master: standby] > en
host [master: standby] # conf t
host [master: standby] (config) # sho ver
```

The output may resemble to the following:

```
Product name:      Content Analytics
Product release:   4.0.rc2
```

```
Build ID:          #1-dev
Build date:        2017-04-20 05:20:28
```

Note: When upgrading the image on nodes in a cluster, execute the following command on the master node. This is required so that role of the master node is the same as it was before the upgrade.

```
host [master: standby] (config) # cluster master self
```

Restart Processes on GMS Nodes

1. Execute the following commands on both the GMS nodes:

```
> en
# conf t
(config)# pm process gms_server launch auto
(config)# pm process pgsqld launch auto
(config)# wr mem
(config)# pm process pgsqld restart
(config)# pm process gms_server restart
```

Upgrade MURAL and Bulkstats Databases on Insta Cluster

After the patch is installed successfully, wait for 15 minutes before upgrading the MURAL and bulkstats databases on the Insta cluster.

1. Perform the following steps to upgrade the schema for new XML for bulkstats cubes from the master iNSTA node:

   a. Execute the following commands:

   ```
   > en
   # _shell
   # cli -m config
   (config)# insta infinidb upgrade
   (config) # quit
   ```

   b. Wait for 10 to 15 minutes for the infinidb adaptor to come up, and then execute the following commands:

```
# echo "insta infinidb get-status-info" | cli -m config |
grep "instance 0 service status :"
```

Ensure that the output contains the following: `XML_UPGRADE_WAIT_FOR_ACTION`.

c. Execute the following command to upgrade the MURAL database:

```
# cli -m config
(config)# insta instance 0 cube-schema-upgrade summary
```

When the new dimension "apn" is displayed as added to cubes, execute the following command:

```
(config)# insta instance 0 cube-schema-upgrade apply
```

d. Wait for 10 minutes, and then execute the following commands to upgrade the Bulkstats database:

```
(config)# insta instance 1 cube-schema-upgrade summary
(config)# insta instance 1 cube-schema-upgrade apply
(config) # quit
```

e. Wait for 5 minutes, then execute the following command:

```
# echo "insta infinidb get-status-info" | cli -m config |
grep -A3 "instance 0 service status :"

```

Ensure that the output resembles as follows:

```
Insta instance 0 service status : RUNNING
Infinidb Adaptor status : Adaptor Running
```

f. Execute the following command to restart iNSTA process:

```
# cli -m config
(config)# pm process insta restart
(config)# quit
```

2. Perform the following steps to restart insta on the standby node:

   a. Execute the following commands:

   ```
   > en
   # _shell
   # cli -m config
   (config)# insta instance 1 cube-schema-upgrade apply
   (config)# insta instance 0 cube-schema-upgrade apply
   (config)# pm process insta restart
   ```

   b. Wait for a few minutes and check the status of infinidb by executing the following command:

   ```
   (config)# insta infinidb get-status-info
   (config)# quit
   ```

   The output may resemble as follows:

   ```
   Infinidb Install status : INSTALLED
   insta instance 0 service status : RUNNING
   Insta instance 1 service status : RUNNING
   Insta instance 2 service status : RUNNING
   ```

3. Execute the following command on the master iNSTA node to check the installation status:

   ```
   # cli -t "en" "conf t" "insta infinidb get-status-info" | grep
   "Insta instance 0 service status"
   ```

   The output may resemble as follows:

   ```
   Insta instance 0 service status : RUNNING
   ```

   Execute the following command:

   ```
   # cli -t "en" "conf t" "insta infinidb get-status-info" | grep
   "Infinidb Adaptor status"
   ```

   The output may resemble as follows:

```
Infinidb Adaptor status : Adaptor Running
```

Execute the following command:

```
# cli -t "en" "conf t" "insta infinidb get-status-info" | grep
"Active insta nodes"
```

The output may resemble as follows:

```
Active insta nodes : <comma separated list of insta IPs>
```

Note the status as RUNNING and list of active iNSTA nodes in the output.

4. Execute the following command on both the iNSTA nodes (master and standby) to ensure that 55555 port is in listen mode:

   ```
   # netstat -napto | grep -i 55555
   ```

   The output may resemble to the following sample:

   ```
   tcp        0      0 :::55555                    :::*
                LISTEN      2730/insta          off (0.00/0/0)
   ```

5. Execute the following command:

   ```
   # idbmysql database_mural -e "select stringId from stringid_
   apn_group_idmap where stringName = 'Unidentified';"
   ```

   APN string ID is displayed in the output of this command. Note this ID.

   The output may resemble to the following sample:

   ```
   +----------+
   | stringId |
   +----------+
   |  1000005 |
   +----------+
   ```

   Run the following script:

   ```
   # python /opt/etc/scripts/apnCubeMigration.py -d database_
   mural -a <id>
   ```

Here, `<id>` is the APN string ID that is obtained from the preceding command. For example, for the preceding output run the script as follows:

```
# python /opt/etc/scripts/apnCubeMigration.py -d database_mural -a 1000005
```

Wait for the script to run completely.

Post-Upgrade Tasks

After upgrading the image and applying required patches, perform the tasks described in this chapter:

- "Activate the New XML from the Master GMS Node" below
- "Apply Configurations to Nodes" on page 24
- "Restart Processes" on page 36

Activate the New XML from the Master GMS Node

1. (Mandatory for Service ID Report and Data Warehouse Integration (DWI) features) Execute the following command to add templates that are not available in ciscoConfig file:

   For Service ID Report:

   ```
   > en
   # _shell
   # cd /config/gms/Profiles/Custom/
   # ln -fs /opt/deployment/GMS_
   Templates/oozie/App/serviceId/workflow_serviceId_apn_with_
   timeout.xml
   ```

   For DWI:

   ```
   > en
   # _shell
   # cd /config/gms/Profiles/Custom/
   # ln -fs /opt/deployment/GMS_Templates/oozie/App/DWH/workflow_
   dwh_with_timeout_jobs.xml
   ```

2. Perform the following steps to run MMI for generating new XML:

   a. Identify the manually updated values in cli. You will have to update these manually after the updating GMS. You must also update the collector related changes before the XML is activated.

The following attributes must be noted:

```
collector modify-instance 1 modify-adaptor edrflow
modify-file-if flowFile filename-format

collector modify-instance 1 modify-adaptor edrflow
modify-file-if flowFile input-directory

collector modify-instance 1 modify-adaptor edrflow
modify-file-if flowFile backup-file-expiry-period

collector modify-instance 1 modify-adaptor edrflow
modify-file-if flowFile backup-directory

collector modify-instance 1 modify-adaptor edrhttp
modify-file-if httpFile filename-format

collector modify-instance 1 modify-adaptor edrhttp
modify-file-if httpFile input-directory

collector modify-instance 1 modify-adaptor edrhttp
modify-file-if httpFile backup-file-expiry-period

collector modify-instance 1 modify-adaptor edrhttp
modify-file-if httpFile backup-directory
```

The following properties, which are configured through mmi, can be updated during mmi:

- `adaptor.<>.numthreads`— can be obtained by searching for the string "num_threads" in the current running configuration.
- `compute.dfs.datanode.data.dir`— can be obtained by searching for the string "dfs.datanode.data.dir" in the current running configuration.

- `rubix sizing properties`— can be obtained by searching for the string "tomcatInstanceMaxSize, initialJavaHeapSize, rubixFQDN, mailSupport, mailNotificationSender, mailNotificationSender".

b. Execute the following command:

```
# sh /opt/etc/scripts/mmi/cisco_mural_deployment_
wizard.sh
```

Use the existing active.xml as input XML to retain the hardware configuration. Note the name of the generated XML provided on the first page.

To use HIVE, perform the following steps:

a. Edit the file, `/opt/deployment/GMS_Templates/hive/hive_mural.xml` as follows:

- Locate the property name, hive.metastore-ip.
- In the `${APPLICATION.PostgreSQL.|pgInstanceID|.Mgmt.VIP}` property value, replace `pgInstanceId` with 0 if the Starter or Medium pack is used. Specify 1 if the Standard pack is used. For example, it is `${APPLICATION.PostgreSQL.0.Mgmt.VIP}` for Starter pack.

b. Execute the following commands to create soft link for `/opt/deployment/GMS_Templates/hive/hive_mural.xml` file at the location, `/config/gms/Profiles/Custom/`:

```
# cd /config/gms/Profiles/Custom/
# ln -fs /opt/deployment/GMS_Templates/hive/hive_
mural.xml
```

c. Execute the following commands to create soft link for `/opt/deployment/GMS_Templates/oozie/Feed/edrhttp_edrflow/workflow_hive_edrhttp_edrflow_with_timeout.xml` at the location `/config/gms/Profiles/Custom/`:

```
# cd /config/gms/Profiles/Custom/
# ln -fs /opt/deployment/GMS_
Templates/oozie/Feed/edrhttp_edrflow/workflow_hive_
edrhttp_edrflow_with_timeout.xml
```

3. Launch the GMS UI using the URL, https://<IP of GMS server>/configure.

4. Open the new XML from the GMS UI and perform the following tasks:

   a. On the **Application** tab, update the Application Name "workflow" instance "0".

   b. (Mandatory for Service Id Report and DWI features) Add application profile from the list: "workflow_custom_dwh_with_timeout_jobs.xml" and "workflow_custom_serviceId_apn_with_timeout.xml".

   c. Attach workflow_custom_hive_edrhttp_edrflow_with_timeout.xml template on workflow application on the NN cluster.

   d. Remove the old hive_mural.xml and attach the new hive_custom_ mural.xml template on hive application on the NN cluster.

   e. Save the XML after validation.

5. Manually disable auto bin sliding on Collector in the final XML. Change the Property_Type tag value to "disable".

```
<Property>
              <Property_Name>
adaptor.bulkStats.binning.autoBinSlide</Property_Name>
              <Property_Value>disable</Property_Value>
              <Property_Type>default</Property_Type>
              <Property_DataType>string</Property_DataType>
              <Property_OriginalName>adaptor.${adaptor-
name}.binning.autoBinSlide</Property_OriginalName>
       </Property>
       <Property>
              <Property_Name>
```

```
adaptor.edrflow.binning.autoBinSlide</Property_Name>
              <Property_Value>disable</Property_Value>
              <Property_Type>default</Property_Type>
              <Property_DataType>string</Property_DataType>
              <Property_OriginalName>adaptor.${adaptor-
name}.binning.autoBinSlide</Property_OriginalName>
       </Property>
       <Property>
              <Property_Name>
adaptor.edrhttp.binning.autoBinSlide</Property_Name>
              <Property_Value>disable</Property_Value>
              <Property_Type>default</Property_Type>
              <Property_DataType>string</Property_DataType>
              <Property_OriginalName>adaptor.${adaptor-
name}.binning.autoBinSlide</Property_OriginalName>
       </Property>
```

6. Execute the following commands to copy data from standby collector to master node:

```
> en
# _shell
# echo show running-config full | cli -m config | grep
"flowFile input-directory" | awk '{print $9}'
# echo show running-config full | cli -m config | grep
"httpFile input-directory" | awk '{print $9}'
# echo show running-config full | cli -m config | grep
"bulkStatsFile" |grep  "input-directory" | awk '{print $9}'
# scp <dir>/* admin@<collector-vip>:/<dir>
```

Here, `dir` contains the output of these commands.

7. From the master GMS node, activate the XML that has only Collector related changes in templates:

```
> en
# _shell
# cli -m config
(config)# gms config <New-XML> activate
```

Note: Allow the runtime activation to complete so that it applies the new template changes. This process may take 5 to 7 minutes as the activation process applies the new changes as per the new activated XML and will only stop and start Collector Process on all 4 Collector nodes including the NameNodes. Execute the following commands to verify the status:

```
> en
# _shell
# cli -m config
(config)# gms config-runtime show status all
```

If any error is encountered, correct the error and execute the following command:

```
(config)# gms config-runtime cluster cluster-name
<clustername> all
```

Here, `<clustername>` refers to the cluster to which the erroneous nodes belong. If the erroneous nodes belong to different clusters, execute thepreceding command on all those clusters.

Apply Configurations to Nodes

- Log in to the master iNSTA node and execute the following commands to update UI metadata:

```
> en
# _shell
# psql -U postgres -d rubixdb < /opt/etc/scripts/anomaly_
migration.sql
```

- (Mandatory for Service ID Report) Log in to both NameNodes and execute the following commands to enable service ID in flow feed file:

```
> en
# _shell
# python /opt/etc/oozie/SolutionConfigs/EnableServiceId.py yes
```

- (Mandatory for DWI) Log in to both NameNodes and execute the following commands:

```
# cli -m config
(config)# internal set modify -
/tps/process/oozie/jobs/TransferJob/actions/TransferOozieActio
n/attribute/destNamenode/value value string <Dest cluster
name>
```

To enable webhdfs on both the clusters, ensure that each datanode and namenode has the source and destination IP addresses mapped in the /etc/hosts file. For example, the hosts.txt file may resemble as following:

```
192.168.194.121      machine-194-121<dest NN node>
192.168.194.122      machine-194-122<dest DN node>
192.168.192.148      machine-192-148<src NN>
192.168.192.149      machine-192-149<dest DN>
192.168.192.151      machine-192-151<dest DN>
192.168.192.152      machine-192-152<dest DN>
192.168.192.41      <Dest cluster name>
```

To change the required IP filter rules on all the nodes of source as well as destination cluster, identify the INPUT rule to be removed:

Execute the following command:

```
# echo "show ip filter configured" | cli -m config | grep DROP
| grep 50070
```

The output may resemble as follows:

```
 71   DROP    tcp     all                    all
dpt 50070
```

Execute the following command:

```
# echo "show ip filter configured" | cli -m config | grep DROP
| grep 50075
```

The output may resemble as follows:

```
 75  DROP    tcp     all                  all
dpt 50075
```

The first column in the preceding output provides the rule number. Execute the following command (as an example) to remove the rule:

```
# cli -m config -t "no ip filter chain INPUT rule 75"
# cli -m config -t "no ip filter chain INPUT rule 71"
```

- Perform the following steps on both the name nodes (as specified):

  a. Execute the following commands to verify the difference between the current and earlier configuration:

  ```
  > en
  # _shell
  # cli -t "en" "conf t" "show running full" >
  /data/Backup_files/`hostname`_configCurrent.txt
  ```

  b. Verify the difference between the current (/data/Backup_files/*_configCurrent.txt) and previous ((/data/Backup_files/* _configBackup.txt)) configuration files.

  Execute the commands that are found.

  Note: Following are the expected commands to be executed on oozie subshell due to feature updates. Take a note of default values corresponding to the properties as these commands will overwrite the values.

  ```
  > en
  # config t
  (config)# internal set modify -
  ```

```
/tps/process/oozie/dataset/
OutboundRoamingReportOutputDataset/attribute/startTime/va
lue value string 2013-01-01T00:00Z
(config)#  internal set modify -
/tps/process/oozie/jobs/OutboundRoamingReport/attribute/j
obStart/value value string 2017-05-02T04:00Z
```

Execute the following commands to set the start time of Hive jobs (AnomalyHiveHourly, AnomalyHiveDaily, AnomalyHiveAggDaily, BSAgg15min, bsHiveJob, and AnomalyHiveMonthly). Following is an example for the AnomalyHiveMonthly job:

```
(config)# _shell
# hdfs dfs -cat /data/AnomalyHiveMonthly/done.txt
```

The output may resemble as follows:

```
2016-07-28T00:00Z
```

Note the timestamp in the preceding output and execute the following commands:

```
# cli -m config
(config)# internal set modify -
/tps/process/oozie/jobs/AnomalyHiveMonthly/attribute/jobS
tart/value value string 2016-07-28T00:00Z
```

Repeat this step for all the Hive jobs.

Execute the following command to set the time zone of jobs:

```
# cat /data/Backup_files/`hostname`_configBackup.txt |
grep 'jobTimezone/v'  |egrep
'AnomalyDailyAggregation|AnomalyHiveAggDaily|AnomalyHiveD
aily'
```

The output may resemble as follows:

```
internal set modify -
```

```
/tps/process/oozie/jobs/AnomalyDailyAggregation/attribute
/jobTimezone/value value string UTC
internal set modify -
/tps/process/oozie/jobs/AnomalyHiveAggDaily/attribute/job
Timezone/value value string UTC
internal set modify -
/tps/process/oozie/jobs/AnomalyHiveDaily/attribute/jobTim
ezone/value value string UTC
```

Note the commands displayed in the outout and execute these as follows:

```
> en
# cli -m config
(config)# internal set modify -
/tps/process/oozie/jobs/AnomalyDailyAggregation/attribute
/jobTimezone/value value string UTC
internal set modify -
/tps/process/oozie/jobs/AnomalyHiveAggDaily/attribute/job
Timezone/value value string UTC
internal set modify -
/tps/process/oozie/jobs/AnomalyHiveDaily/attribute/jobTim
ezone/value value string UTC
```

c. Execute the following command on the master name node:

```
# cli -m config
(config)# pm process tps restart
```

Wait for 10 minutes and execute these commands on standby name node.

Run the following commands to verify Hadoop status:

```
> en
# _shell
# hdfs dfsadmin -report | head -12
```

The output may resemble to the following sample:

```
Configured Capacity: 5624751722496 (5.12 TB)
Present Capacity: 5326934437888 (4.84 TB)
DFS Remaining: 1634512154624 (1.49 TB)
DFS Used: 3692422283264 (3.36 TB)
DFS Used%: 69.32%
Under replicated blocks: 852
Blocks with corrupt replicas: 0
Missing blocks: 0

-------------------------------------------------
Datanodes available: 4 (4 total, 0 dead)
```

Here, `4` denotes the number of Compute nodes available in the setup.

Retain Manual Changes in IBs

Any changes that are done manually in an IB are lost during the upgrade.

Perform the following steps on the master name node to retain the changes made in IBs:

1. Compare the diff between backup IB of inbox and `/opt/catalogue/atlas` of image.
2. If an IB is found to be overwritten or modified after the update all ib is performed, create a map file containing new entries and an empty line at the end. The structure of the new entry must be similar to the existing entries in the map file.

   Note: The IBs, apnGroup.map, ratidtype.map, ipSgsn.map, ipGgsn.map, and dcRegionArea.map are non-empty and contain all required entries.
3. Execute the following commands for bulk update:

```
> en
# _shell
# pmx subshell aggregation_center
```

```
(aggregation_center)# edit ib <map file to be updated> add
bulk <path to bulk file>
```

Update the IBs

Perform the following steps:

1. Execute the following commands to update the DPI IB:

```
> en
# _shell
# pmx subshell aggregation_center
(aggregation_center) # update all ibs from image
(aggregation_center) # generate urlcat ib from image
(aggregation_center) # show urlcat version
(aggregation_center) # generate all ibs
(aggregation_center) # push all ibs
(aggregation_center) # quit
```

2. Execute the following commands to update the Bulkstats IB:

```
# pmx subshell bulkstats
(bulkstats) # update all ibs from image
(bulkstats) # generate all ibs
(bulkstats) # push all ibs
(bulkstats) # quit
```

3. Execute the following commands to update the Anomaly IB:

```
# pmx subshell anomaly
(anomaly) # update all ibs
(anomaly) # quit
```

Adding Home Network for Outbound Roaming Report

You can generate a report to understand the usage of the subscribers who are roaming (outbound roamers). MURAL provides device, content and Service Provider usage reports per roaming partner based on the feature.

For the outbound roaming report feature, you must configure the mobile country code (MCC)/mobile network code (MNC) of the required outbound roamers.

Perform the following steps:

1. Verify if the outbound roaming report feature is enabled:

```
> en
# _shell
# grep -e "emitCubes" /opt/etc/oozie/EDR/app/Annotations.json
"emitCubes":"true",
```

   Note: The value of `"emitCubes"` as `"true"` signifies that the feature is enabled.

2. Execute the following commands to add the MCC/MNC of the outbound roamers:

```
> en
# conf t
(config)# pmx subshell aggregation_center
pm extension (aggregation center)> edit ib homeNetwork.list
add
mcc: <mobile country code of roamer>
mnc: <mobile network code of roamer>
```

   Add all the possible MCC/MNC combinations. To view all the added values:

```
pm extension (aggregation center)> show ib homeNetwork.list
1 [440][001]
```

3. Execute the following commands to add the IB to the system:

```
pm extension (aggregation center)> push ib homeNetwork.list
```

   Note: Execute the code, `opt/etc/scripts/validateHomeNwList.py` before running the core job to check any overlapping MNC prefixes for a given MCC in the home network list.

   For example,

```
# python /opt/etc/scripts/validateHomeNwList.py
```

The output may resemble to the following sample:

```
Fetching information from /data/ib/inbox/homeNetwork.list
Fetching information from /data/ib/inbox/globalNetwork.list
For 440:002 also exits 440:00
For 310:02 also exits 310:020
```

Clear the Disk Cache for the Atlas Application

Execute the following commands to clear the existing disk cache for the atlas application:

```
> en
# _shell
# rm -rf /data/diskstore/UI-EDR
# mkdir -p /data/diskstore/UI-EDR
# touch /data/diskstore/UI-EDR/disk_mounted
```

Configure System Monitoring Interface

System Monitoring Interface (SMI) is a dashboard that can be used to display MURAL's system-related health information for the rolling 24-hours period.

Benefits of SMI:

- The system administrator can check the health of the MURAL system and contact Technical Support if a condition is not being met.
- The system administrator can use the dashboard to provide assurance to customer that MURAL services are working fine as per the benchmarking.

Note: You can install SMI only in a Standard pack or a Medium pack.

To configure SMI, perform the following steps:

1. Log into the RGE node master node.
2. Execute the following command to copy the active GMS file from the GMS

node to the RGE master node:

```
> en
# _shell
# cd /var/home/root
# scp admin@<GMS Master IP/Hostname>:/config/gms/active-
configuration/resolved_active.xml .
```

3. Execute the following commands to install SMI on the RGE master node:

```
# cd /var/home/root
# tar -xvzf pkg_dashboard.tgz
#  ./install_dashboard.sh
```

The output may resemble as follows:

```
    ----------------------------------------------------------
---------------------
    install_dashboard.sh:msg: Usage.......
    ----------------------------------------------------------
---------------------
    install_dashboard.sh:input: Enter TAR Filename(with full
path) : /var/home/root/dashboard.tgz
    install_dashboard.sh:input: Enter GMS XML Filename(with
full path) : /var/home/root/resolved_active.xml
    install_dashboard.sh:input: Enter Management Network Name
used in above XML[MgmtNW]: MgmtNW
    install_dashboard.sh:msg: Network Name is : MgmtNW

    ----------------------------------------------------------
---------------------


    install_dashboard.sh:msg: Extracting tar ball into node
...
    .
    .
    .
    .
    .

    Stopping pm:
[  OK  ]

    Starting pm:
[  OK  ]

    install_dashboard.sh:msg: Dashboard configuration is done
successfully!!!!

    install_dashboard.sh:msg: Start Dashboard from UI by
http://<This node IP>:15443

    #
```

Note: Provide the value for "Management Network Name" attribute as specified in the resolved_active.xml file.

4. Execute the following commands to check if SMI processes have started successfully:

   - To check if grafana is up:

     ```
     # ps -ef|grep grafana
     ```

     The output may resemble as follows:

     ```
     admin     4300  2653  0 14:43 ?        00:00:00
     /opt/grafana/bin/grafana-server -homepath=/opt/grafana
     ```

```
admin     6669 31484  0 14:44 pts/1    00:00:00 grep
grafana
```

- To check if graphite and carbon-cache are up:

```
# ps -ef|grep graph
```

The output may resemble as follows:

```

admin     3313  2653  0 14:43 ?        00:00:00
/usr/local/bin/python /opt/graphite/bin/carbon-cache.py -
-nodaemon --logdir=/var/log --pidfile=/var/run/carbon-
cache.pid start

admin     5546  2653  1 14:43 ?        00:00:00
/usr/local/bin/python /usr/local/bin/gunicorn_django --
bind=127.0.0.1:12040 --log-file=/var/log/gunicorn-
graphite/gunicorn.log --preload --
pythonpath=/opt/graphite/webapp/graphite --
settings=settings --workers=3 --pid=/var/run/gunicorn-
graphite/gunicorn-graphite.pid

admin     5753  5546  0 14:43 ?        00:00:00
/usr/local/bin/python /usr/local/bin/gunicorn_django --
bind=127.0.0.1:12040 --log-file=/var/log/gunicorn-
graphite/gunicorn.log --preload --
pythonpath=/opt/graphite/webapp/graphite --
settings=settings --workers=3 --pid=/var/run/gunicorn-
graphite/gunicorn-graphite.pid

admin     5773  5546  0 14:43 ?        00:00:00
/usr/local/bin/python /usr/local/bin/gunicorn_django --
bind=127.0.0.1:12040 --log-file=/var/log/gunicorn-
graphite/gunicorn.log --preload --
```

```
pythonpath=/opt/graphite/webapp/graphite --
settings=settings --workers=3 --pid=/var/run/gunicorn-
graphite/gunicorn-graphite.pid

admin     5779  5546  0 14:43 ?        00:00:00
/usr/local/bin/python /usr/local/bin/gunicorn_django --
bind=127.0.0.1:12040 --log-file=/var/log/gunicorn-
graphite/gunicorn.log --preload --
pythonpath=/opt/graphite/webapp/graphite --
settings=settings --workers=3 --pid=/var/run/gunicorn-
graphite/gunicorn-graphite.pid

admin     6702 31484  0 14:44 pts/1    00:00:00 grep
graph
```

Note the running statuses of the graphite and carbon-cache processes as highlighted in the preceding sample output.

5. Execute the following commands:

```
# cli -m config
(config)# cluster master self
(config)# quit
```

Restart Processes

1. Run the following commands on all the collector nodes to restore the collector bulkstats configuration from configuration backup:

```
> en
# _shell
# echo -e "\nen\n\nconf t\n\n"  > /data/Backup_files/coll_
bsConfigRestore.cli
# grep "add-file-if bulkStatsFile" /data/Backup_
files/`hostname`_configCurrent.txt | sed 's/add/delete/' >>
/data/Backup_files/coll_bsConfigRestore.cli
```

```
# cat /data/Backup_files/coll_bsGW_config.txt >> /data/Backup_
files/coll_bsConfigRestore.cli
# cli -x -h /data/Backup_files/coll_bsConfigRestore.cli
```

For more information about gateway configurations, see "Take a Backup of Gateway Configurations" on page 8

2. Perform the following steps to configure SSH session limit:

   a. Execute the following commands to set the concurrent SSH limit to 50:

   ```
   > en
   # conf t
   (config) # ssh server max-startups 50
   (config) # ssh server max-sessions 50
   (config) # write memory
   (config) # pm process sshd restart
   ```

   b. Repeat steps 1 through 2 on the standby Collector node.

3. Log into Collector nodes and restart collector:

```
> en
# conf t
# pm process collector restart
```

4. Log into corresponding UI nodes, and execute the following commands:

```
> en
# _shell
# /bin/rm -f /data/instances/*/1/bin/blazed*
# cli -m config
(config)# no pm process rubix launch auto
(config)# pm process rubix restart
```

```
(config) # rubix modify-app <app-name> modify-instance 1
enable
```

```
(config) # rubix modify-app <app-name> enable
(config) # write memory
```

Where, <app-name> is as follows. Start the applications in the same order as mentioned below.

- atlas
- bulkstats
- reportAtlas
- rge
- ruleEngine

5. Update the `/etc/host` mapping as follows:
   1. Map the UI node to domain. For example, <VIP of UI> domain.
   2. Map the RGE and other rubix application nodes to domain-report. For example, <VIP of RGE cluster> domain-report.
6. Access the UI by going to the URL: https://domain-name:6443/ through your browser.

   The domain name to be used is the one which was provided at the time of initial configuration through GMS for the UI nodes configuration details.

   For example: URL: https://demo.cisco.com:6443/
7. Perform the following steps to access SMI user interface:
   1. Log into MURAL with administrator credentials.
   2. On the top-right corner, click ⠿ > **System Monitoring Interface**.
8. Log into master NameNode and run all jobs:

```
> en
# _shell
# pmx subshell oozie run job all
```

Enable the Collector in the Live Mode

- Execute the following commands to enable the collector again in the live mode:

```
> en
# _shell
# cli -m config
(config)# collector modify-instance 1 modify-adaptor bulkStats
auto-bin-slide enable
(config)# collector modify-instance 1 modify-adaptor edrflow
auto-bin-slide enable
(config)# collector modify-instance 1 modify-adaptor edrhttp
auto-bin-slide enable
(config)# wr me
(config)# pm process collector restart
```