

Cisco MURAL Software Installation Guide

Version 3.2 MR2

Last Updated March 4, 2014

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.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

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.

Cisco MURAL Software Installation Guide

Copyright © 2014, Cisco Systems, Inc.All rights reserved.

Table of Contents

Table of Contents **5**

Installation Overview **8**

Prerequisites 8

System Components 8

Hardware 10

Installation Package 10

Customer Information Questionnaire (CIQ) 11

Installation Process Steps 11

Verifying UCS Hardware Configuration for MURAL **13**

Preparing for Initial Configuration 13

Hardware Topology 13

Installing the Blades 14

Connecting UCS SAN Uplinks 19

Connecting UCS Network Uplinks 20

Setting the Base Configuration for the UCS System 21

Configuring UCS for MURAL **23**

Configuring Direct Attachments to External SANs 23

System Profile Settings for the Cisco UCS Manager 40

Manufacturing the GMS Blade **45**

Setting Up the GMS Node **53**

Defining the Deployment Topology Using GMS **55**

Accessing the GMS Interface 55

Loading the Base Configuration File 56

Verifying Server Details for Each Node 57
Verifying Global Settings 60
Configuring the Node Clusters 62
Specifying Component Profiles and Application-Specific Settings 63
Validating the Configuration 73
Manufacturing All Other Blades 75
Setting Up Data Storage On the EMC 80
EMC Hardware Installation Prerequisites 80
Configuring Your Base IP to Access the EMC System 81
Registering All MURAL Nodes on the EMC 81
Creating and Assigning RAID Groups and LUNs 85
Creating Storage Groups 89
Adjusting Caching on the EMC 93
Return to GMS and Configure the LUN WWIDs 95
Installing MURAL on the UCS Nodes 96
Troubleshooting Node Installation 97
Tuning the System for Performance 97
Verify the Processes are Running 98
Generate and Push the Information Bases 100
Tethering 100
Configuring IBs for EDR 102
Configuring DCs and Gateways For All IBs 104
Synchronize the EDR and BulkStats IBs on the Standby Collector Node 110
Fetching BulkStats IBs 113

 Copyright © 2014, Cisco Systems, Inc.

Create File for Uncategorized URL, UA, and TAC Reports 113

Single Certificate Installation **115**

Backing Up and Generating the Keystore Files 115

Downloading the Signed Certificate 117

Downloading the CA Certificate 117

Installing the Signed Certificate in the Keystore 118

Processing the Data **121**

Setting Up a New User for the ASR in the Collectors 121

Ingesting the Data Into the System 122

Validating the System Installations **123**

Data Validation on the Collector Nodes 123

Data Validation on Compute Blades (Data Nodes) 125

Data Validation on Insta Blades 125

Validate Bulk Stats Data on the Insta Blade 127

Start UI Processes and Verify Data 128

Mandatory Parameters for Incoming ASR Files **131**

Mandatory Attributes for Flow EDRs for MURAL 131

Mandatory HTTP EDR Attributes for MURAL 132

ASR-Side Configuration 132

Glossary **135**

Installation Overview

This document describes how to install the Cisco Mobility Unified Reporting and Analytics (MURAL) application. MURAL provides Web-based reporting and analytics abilities for deep packet inspection (DPI) data emerging from your network.

Prerequisites

Before beginning the Cisco MURAL installation, you should have a working knowledge of the following:

- Linux
- Cisco UCS 2.1
- EMC storage devices

Prior to installing the application, we recommend that you:

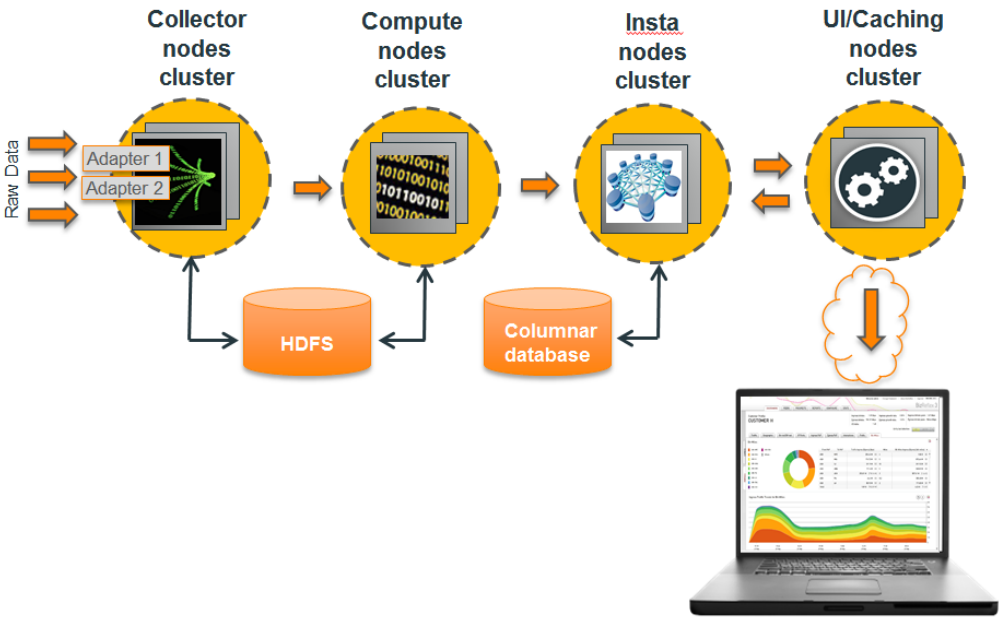
- Review the MURAL 3.2MR2 (or later) Release Notes
- Complete a training course on MURAL
- Have an understanding of Cisco UCS Server Blade hardware administration
- Locate your Customer Information Questionnaire (CIQ). See "Customer Information Questionnaire (CIQ)" on page 11
- Verify that each Fabric Interconnect is connected to the SAN controllers through fibre cables.
- Verify all UCS B200 M2/M3 blade servers have been installed physically on the UCS 5108 Blade Server Chassis and connected to the UCS 6248 Fabric Interconnect for each of the types of nodes (Collector, Compute, Insta, UI/Caching, and GMS nodes). The number of each type of node is customized for your deployment,

Note: Set up hardware as specified in the bill of materials (BOM).

System Components

The following figure shows the components of the MURAL platform, focusing on how the data flows through the system:

 Copyright © 2014, Cisco Systems, Inc.



The MURAL platform consists of the following nodes, each hosted on blades in the UCS Chassis:

- **Collector node**—Collects data streams pushed to the Reflex platform, interprets the exported flows, enriches them with static data, and assembles data sets. The Collector stores the raw data in the Hadoop file system (HDFS) and sends it to the Compute node. The Collector node cluster can have any number of servers, in pairs for master and standby and uses 1+1 redundancy (transparent failover between pairs of active-active nodes).
- **Compute node**—Analyzes and aggregates the data, creating *data cubes*. The term data cube is a convenient shorthand for the data structure, which is actually multi-dimensional and not limited to the three dimensions of a cube. The Compute node cluster can have any number of servers, depending on your deployment, and uses N+1 redundancy.
- **Insta node**—Stores and manages the processed data in a columnar database. Manages the Insta database, which stores processed data cubes, commonly three to six years' worth. The Insta node cluster has two servers with 1+1 redundancy.
- **UI/Caching (Rubix) node**—Hosts the Rubix engine and data cache. The Rubix engine queries the Insta nodes constantly and when new data is available, it fetches it to store in the data cache, so that it can respond more

quickly to requests from the UI engine. The Caching node is sometimes called the Rubix node. The Caching node uses N+1 redundancy in active-active mode.

- **General Management System (GMS)**—Provides centralized management of the Reflex platform nodes, such as remote manufacturing of blades (installing the MURAL operating system), patch management, monitoring of all nodes and operations, and importing and running node configurations.

Hardware

The MURAL application is hosted on a redundant pair of UCS 5108 Blade Server Chassis. Each chassis includes blades that host the nodes listed above (Collector, Compute, Insta, UI/Caching, and GMS). Data storage is hosted on EMC storage devices.

The data flows that feed the MURAL system are pushed by an ASR 5000 or ASR 5500 platform (hereafter referred to simply as an ASR).

Installation Package

The MURAL software installation package contains the following components:

- An ISO image— For example:
  `md5sum /data/mfgcd-guavus-x86_64-20130925-174106.iso`
  For the exact image name and the MD5 checksum for the software image, refer to the release notes for your release.
- The **mural.xml** file, which you will install on the GMS node. This sample configuration file provides configuration settings for your setup, including a logical grouping of the nodes into two chassis and configuration settings for all of the system components. These settings are based on the details of your network that you supplied in the Customer Information Questionaire (CIQ).
- Any software patches that are available for the release. See the Release Notes for the appropriate release for a complete list of patches.
- Management information bases (MIBs)

 Copyright © 2014, Cisco Systems, Inc.

Customer Information Questionnaire (CIQ)

The Customer Information Questionnaire (CIQ) is an Excel spreadsheet containing a site survey that you completed before beginning the installation process.

The CIQ file contains the following worksheets:

- **Contacts**—Identifies site personnel and their responsibilities
- **Space_Power Req**—Specifies space and power requirements
- **IP Survey**—Specifies physical network connections, VLANs, various interfaces, SNMP/SMTP settings, and so forth
- **Network Diagrams**—Shows the system components and how they are connected
- **Connectivity**—Specifies the details for ports and connections
- **Firewall**—Identifies the firewall changes required for connectivity
- **Alarms**—Lists and describes all SNMP traps supported by the application
- **ASR**—Specifies locations for various ASR information bases (IBs) that are required by the application

Installation Process Steps

The MURAL installation process includes the following steps:

Warning: Skipping a task or performing the tasks out of sequence may cause a misconfiguration that results in system failure.

1. Verify that the UCS hardware is correctly set up and configured for the MURAL system. See "Verifying UCS Hardware Configuration for MURAL" on page 13
2. Set basic MURAL settings on the UCS by:
   - Running the ucs-config script to set configuration parameters for the UCS Fabric Interconnect.
   - Configure the Direct-Attach SAN feature.

   See "Configuring UCS for MURAL" on page 23.
3. Load the MURAL software onto the GMS node, complete the initial configuration of GMS, define the deployment topology for the MURAL system,

and then load the MURAL image onto the remaining blades:

- "Manufacturing the GMS Blade" on page 1
- "Setting Up the GMS Node" on page 1
- "Defining the Deployment Topology Using GMS" on page 1
- "Manufacturing All Other Blades" on page 1

4. Set up the data storage for the MURAL system:
   - Set up the storage. See "Setting Up Data Storage On the EMC" on page 80.
   - Go back and update GMS with the actual WWIDs for LUNs. See "Defining the Deployment Topology Using GMS" on page 55
5. Complete the installation of the MURAL software on the UCS blades:
   - Apply available patches on applicable nodes. For information about available patches, see the Release Notes for 3.2 MR2 or above.
   - Install MURAL on the nodes. See "Installing MURAL on the UCS Nodes" on page 96
6. Configure environmental parameters specific to your environment:
   - "Generate and Push the Information Bases"
   - "Single Certificate Installation " on page 115
   - "Processing the Data " on page 121
7. Verify that the system is working as expected. See "Validating the System Installations" on page 123.

 Copyright © 2014, Cisco Systems, Inc.

Verifying UCS Hardware Configuration for MURAL

The Cisco UCS 5108 Blade Server hosts the blades (nodes) of the MURAL platform. This topic describes

- UCS hardware configuration details to verify—Slot assignments for blades, fabric interconnections, uplinks for UCS SANs, and network uplinks.
- Initial configuration of the UCS—Setting the admin password, setting up the management port IP address, setting up a cluster for the two fabrics, and specifying the default gateway.

Before you begin, verify that you have all items listed in the bill of materials (BOM).

Preparing for Initial Configuration

Before you begin configuring the fabrics, verify the following physical connections:

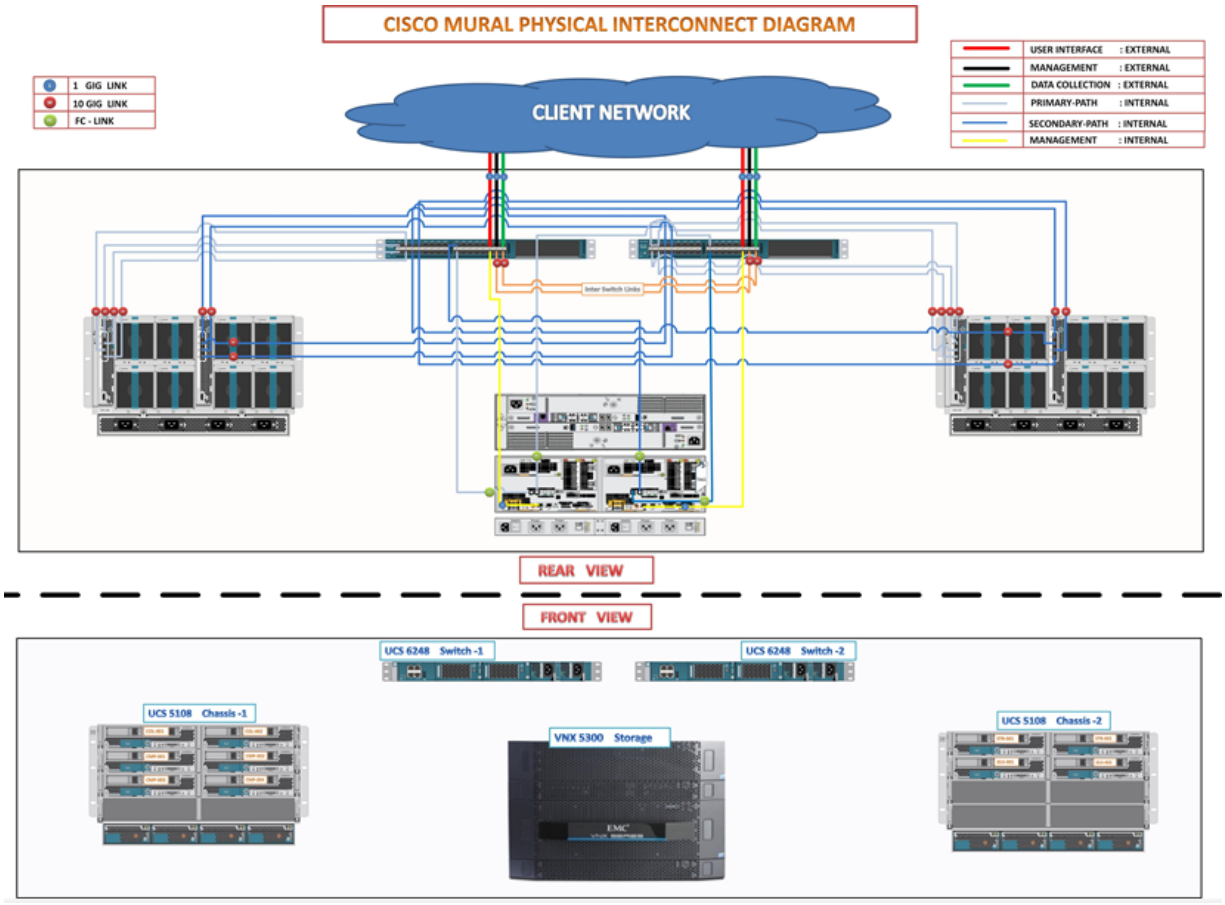
- The Management Ethernet port (mgmt0) is connected to an external hub, switch, or router.
- L1 ports on both fabric interconnects are directly connected to each other.
- L2 ports on both fabric interconnects are directly connected to each other.

Set the console port parameters on the computer terminal (or console server) to:

- 9600 baud
- 8 data bits
- No parity
- 1 stop bit

Hardware Topology

The standard MURAL deployment has been tested and validated for the topology shown in the following figure.



Installing the Blades

Install the blades in each chassis as shown in the following diagram and tables. This is an example only; your configuration might vary. Refer to the CIQ for your setup to make actual slot assignments.

Note: All blades are identical, with the exception of the Rubix/UI blade, which has two to three times more RAM than the other blades.

Slots are numbered 1 through 8 from top to bottom, left to right.

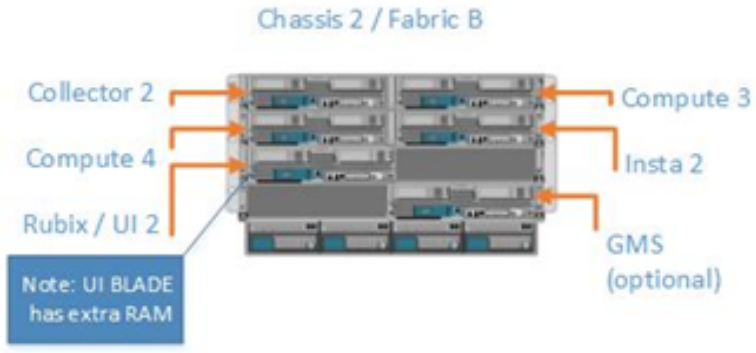
On chassis 1, make connections to Fabric Interconnects as specified in the following table. The slot assignments for Chassis 1 (Fabric A) are shown in the following figure.

Copyright © 2014, Cisco Systems, Inc.



Chassis	Slot	Blade
Chassis 1	Slot 1	Collector 1
Chassis 1	Slot 2	Compute 1
Chassis 1	Slot 3	Compute 2
Chassis 1	Slot 4	Insta 1
Chassis 1	Slot 5	Rubix/UI 1
Chassis 1	Slot 6	GMS 1

On chassis 2, make connections to Fabric Interconnects as specified in the following table. The slot assignments for Chassis 2 (Fabric B) are shown in the following figure.



Chassis	Slot	Blade
Chassis 2	Slot 1	Collector 2
Chassis 2	Slot 2	Compute 3

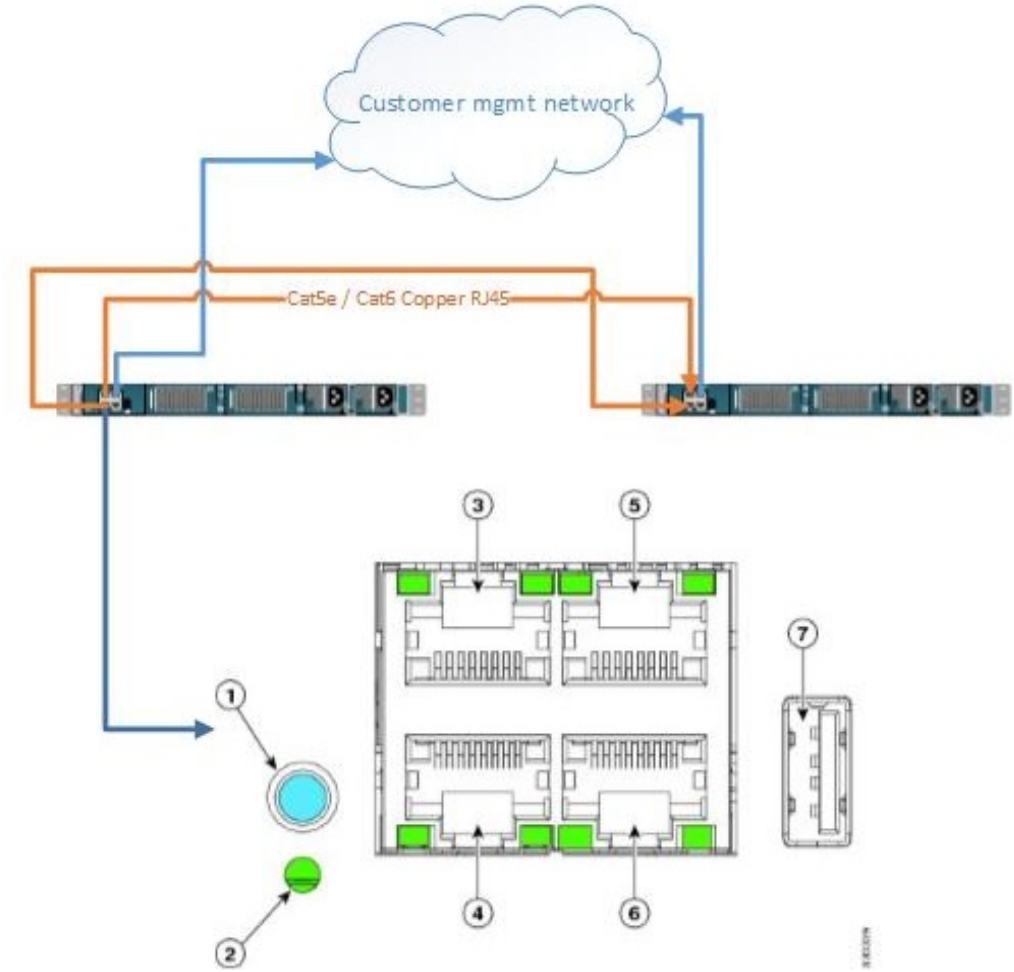
Chassis	Slot	Blade
Chassis 2	Slot 3	Compute 4
Chassis 2	Slot 4	Insta 2
Chassis 2	Slot 5	Rubix/UI 2
Chassis 2	Slot 6	GMS 2 (optional)

Note: When assigning the blades, ensure that **HA Level** is maintained for Collector, Compute, Insta, Rubix (UI/Caching), and optionally GMS nodes. Ensure that HA nodes are not assigned on same chassis.

Physical Connections on the Fabric Interconnects

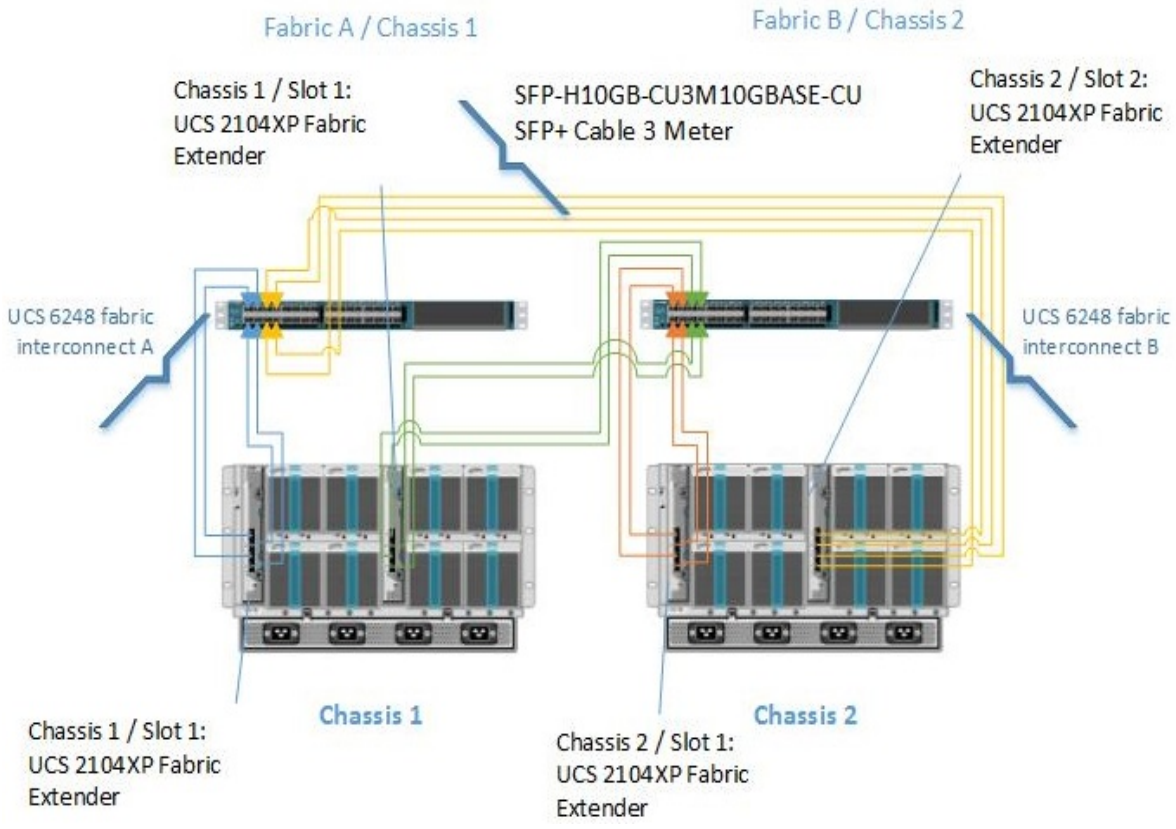
Verify the connections for the Cisco UCS 6248UP 48-Port Fabric Interconnect. The figure below shows a front view and illustrates the connections needed for the MURAL components.

Copyright © 2014, Cisco Systems, Inc.



Number	Description
1	Beaconing LED and Reset button
2	System status LED
3	UCS cross-connect port L1
4	UCS cross-connect port L2
5	Network management port
6	Console port
7	USB port

Make the fabric interconnections as shown in the following figure, which depicts the fabric interconnect wiring:



Chassis 1

Make chassis 1 connections to Fabric Interconnects as shown in the following table:

UCS 2104XP Fabric Extender Slot 1	UCS 6248 Fabric Interconnect A
Chassis 1 : Extender 1 : Port 1	Interconnect A : Port 1
Chassis 1 : Extender 1 : Port 2	Interconnect A : Port 2
Chassis 1 : Extender 1 : Port 3	Interconnect A : Port 3
Chassis 1 : Extender 1 : Port 4	Interconnect A : Port 4
UCS 2104XP Fabric Extender Slot 2	UCS 6248 Fabric Interconnect B

Copyright © 2014, Cisco Systems, Inc.

UCS 2104XP Fabric Extender Slot 1	UCS 6248 Fabric Interconnect A
Chassis 1 : Extender 2: Port 1	Interconnect B: Port 5
Chassis 1 : Extender 2: Port 2	Interconnect B: Port 6
Chassis 1 : Extender 2: Port 3	Interconnect B : Port 7
Chassis 1 : Extender 2: Port 4	Interconnect B: Port 8

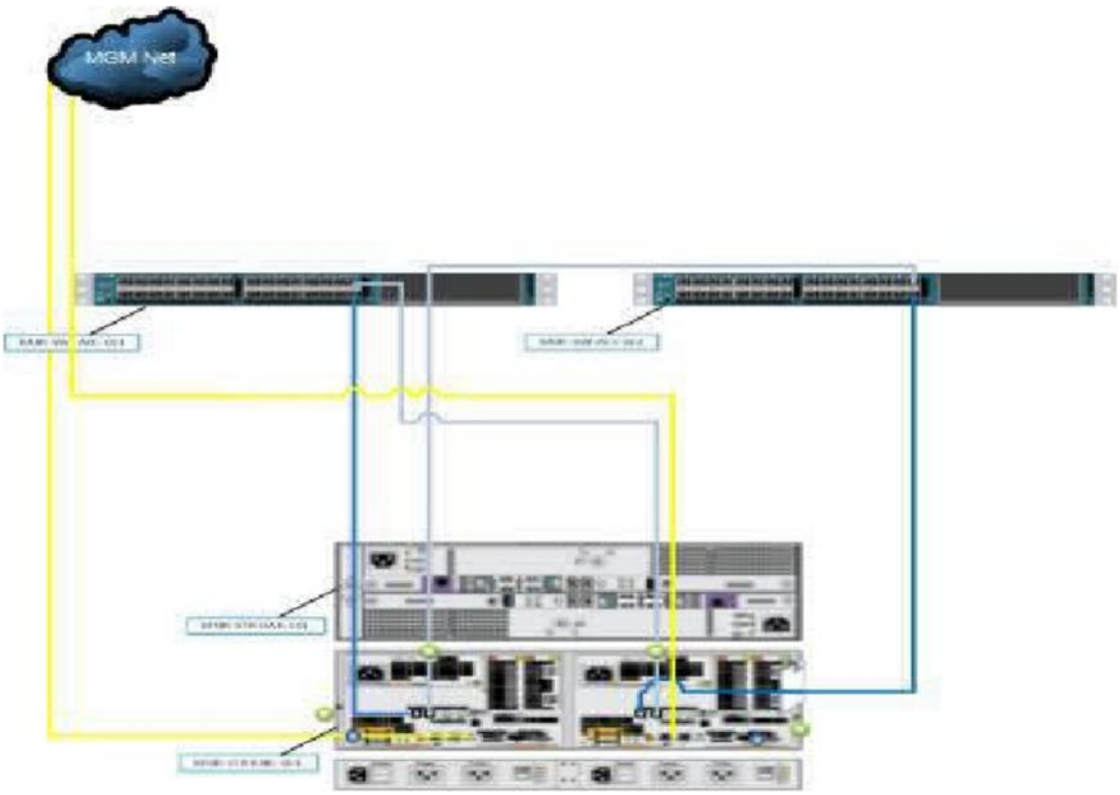
Chassis 2

Make chassis 2 connections to Fabric Interconnects as shown in the following table:

UCS 2104XP Fabric Extender Slot 1	UCS 6248 Fabric Interconnect B
Chassis 2: Extender 1 : Port 1	Interconnect B : Port 1
Chassis 2: Extender 1 : Port 2	Interconnect B: Port 2
Chassis 2 : Extender 1 : Port 3	Interconnect B : Port 3
Chassis 2 : Extender 1 : Port 4	Interconnect B : Port 4
UCS 2104XP Fabric Extender Slot 2	**UCS 6248 Fabric Interconnect A**
Chassis 2 : Extender 2: Port 1	Interconnect A: Port 5
Chassis 2 : Extender 2: Port 2	Interconnect A: Port 6
Chassis 2 : Extender 2: Port 3	Interconnect A : Port 7
Chassis 2 : Extender 2: Port 4	Interconnect A: Port 8

Connecting UCS SAN Uplinks

Connect the UCS SAN uplinks as shown in the following figure and table:

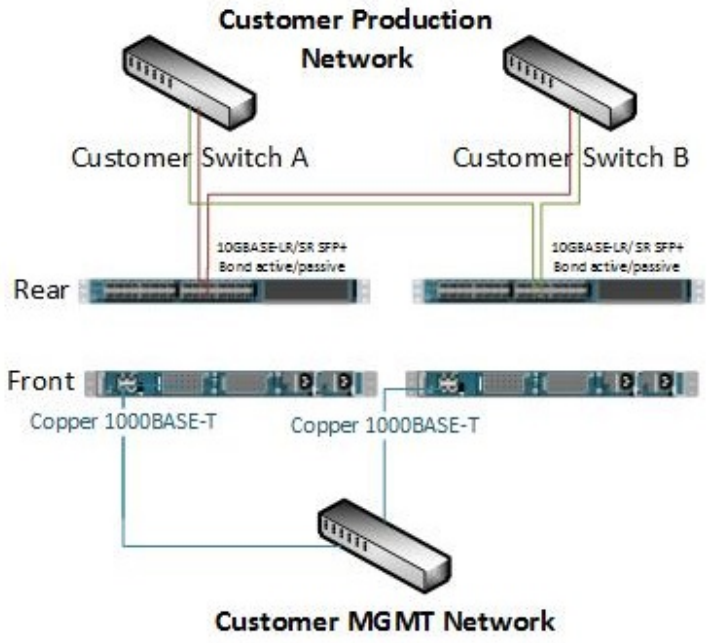


SP-A management port	Customer management switch
SP-B management port	Customer management switch
SP-A FC-A	Fabric A--Port 31
SP-A FC-B	Fabric B--Port 31
SP-B FC-A	Fabric A--Port 32
SP-A FC-B	Fabric B--Port 32

Connecting UCS Network Uplinks

Connect the UCS physical network uplinks as shown in the following figure and table:

Copyright © 2014, Cisco Systems, Inc.



Port	Description
Fabric A--port 17	Customer production network
Fabric B--port 17	Customer production network
Fabric A--port 18	Customer secondary production switch (optional)
Fabric B--port 18	Customer secondary production switch (optional)

Setting the Base Configuration for the UCS System

To set the base configuration for the UCS system and enable the fabrics be brought up of the fabrics, complete the following procedure to set the admin password, set up the management port IP address, set up a cluster for the two fabrics, and specify the default gateway.

To set the base UCS configuration:

1. Connect to the console port of fabric A.

   Note: Use these settings for the console port parameters on the computer terminal or console server—9600 baud, 8 data bits, 1 stop bit, and no parity.

2. For fabric A, set the following parameters:

```
Configuration method: console
Setup mode: setup
New fabric interconnect: Y
Enforce strong password: Y

Admin password: admin-password
Is this Fabric Interconnect part of a cluster:Y
Switch fabric: A
System Name: UCS-name <without a trailing -A or –B>
Mgmt0 IP address: Fab-A-mgmt-port-IP-address>
Mgmt0 Netmask: mgmt-port-IP-netmask
IPv4 default gateway: gateway-address-in-mgmt-subnet
Cluster IPv4 address: Virtual-IP-for-active-node <usually
belonging to the management subnet>
```

   Note: You can also configure the DNS server address and the unit's domain name, but this is not required.

2. Connect to the console port of fabric B, and verify the redundancy cables between the two fabrics are connected. Perform the initial configuration with the following parameters:

```
Configuration method: console
This fabric interconnect will be added to the cluster:Y
Admin password of interconnect: admin-password <used in the
configuration of Fab A>
Mgmt0 IP address: Fab-B-management-port-IP-address
```

You can now log in to the management UI from a web browser at **http://*ip-address-of-cluster***.

 Copyright © 2014, Cisco Systems, Inc.

Configuring UCS for MURAL

To complete the configuration of the UCS for MURAL, run the MURAL configuration scripts for UCS, set up the Cisco UCS Direct-Attached SAN, and set system profile settings for UCS Manager.

A script enables quick configuration for UCS for the MURAL installation. To complete the UCS configuration for MURAL, locate the following files, which can be obtained from either Cisco Advanced Services, or Technical Support:

- **ucs-config-*version-number.txt***(where *version-number* is the most recent version available)—The configuration parameters in this file are used by the UCS configuration script. You must update this file with your local setup details.
- **ucs-config.exp**—Sets configuration parameters for the UCS Fabric Interconnect, servers, LAN, and SAN.

To run the UCS configuration script:

1. Edit the **ucs-config-*version-number*.txt** file—modify each value that is marked with a CIQ label to match the value in your CIQ spreadsheet.
2. Save and rename the completed **ucs-config-*version-number*.txt** file into the same directory as the **ucs-config.exp** script.
3. Verify that you can ping the UCS management IP address:

   /ucs.exp *ucs-mgmt-ip ucs-password*

4. From a Cygwin, Linux, or a Mac terminal, run the script and watch for any errors or issues.

Tip: If the script encounters an error, you can recover by resetting the UCS to **defaults.ssh** from the UCS manager. You need to do this for both A and B sides.

Configuring Direct Attachments to External SANs

This section describes how to set up Cisco UCS Direct-Attached SAN, which enables you to directly-attach a fiber-channel SAN to the Fabric Interconnects.

To configure Direct-Attached SAN, complete these tasks described in the following sections:

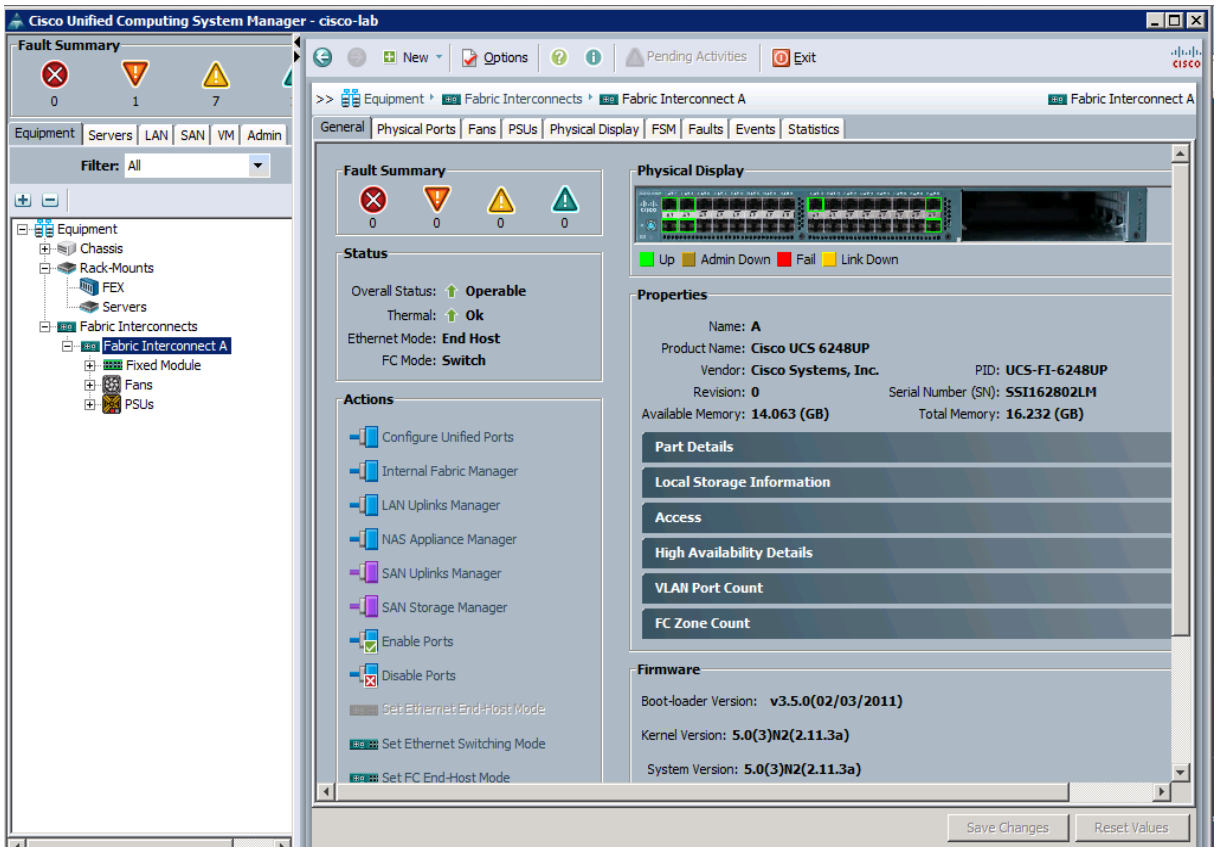
- "Setting Fabric Interconnects to FC Switch Mode" below
- "Creating VSANs for Zoning " on the facing page
- " Designating Storage Ports" on page 27
- "Assigning the Storage Cloud VSAN to the FC Ports" on page 29
- "Confirm Storage Port WWPN is Logged Into the Fabrics" on page 30
- "Creating Storage Connection Policies" on page 33
- "Creating SAN Connectivity Policy" on page 35
- "Configuring SAN Cloud Policy" on page 35
- "Create vHBA Initiator Groups" on page 38
- "Verify Service Profile Templates" on page 39
- "System Profile Settings for the Cisco UCS Manager" on page 40

Setting Fabric Interconnects to FC Switch Mode

To set the Fabric Interconnects into FC Switch Mode:

1. In the UCS Manager, navigate to **Equipment tab** > **Equipment** > **Fabric Interconnect A/B** > **General** tab.
2. Select both **Set FC Switching Mode** and **Set Ethernet Switching Mode**.
3. Click **Save Changes**.
4. If **FC Mode** is not shown as **Switch**, reboot the system.
5. Repeat steps 1 through 3 for the other Fabric Interconnects.

 Copyright © 2014, Cisco Systems, Inc.



Creating VSANs for Zoning

Create one virtual storage area network (VSAN) for each Fabric Interconnect.

Follow these guidelines and recommendations when naming VSANs, including storage VSANs:

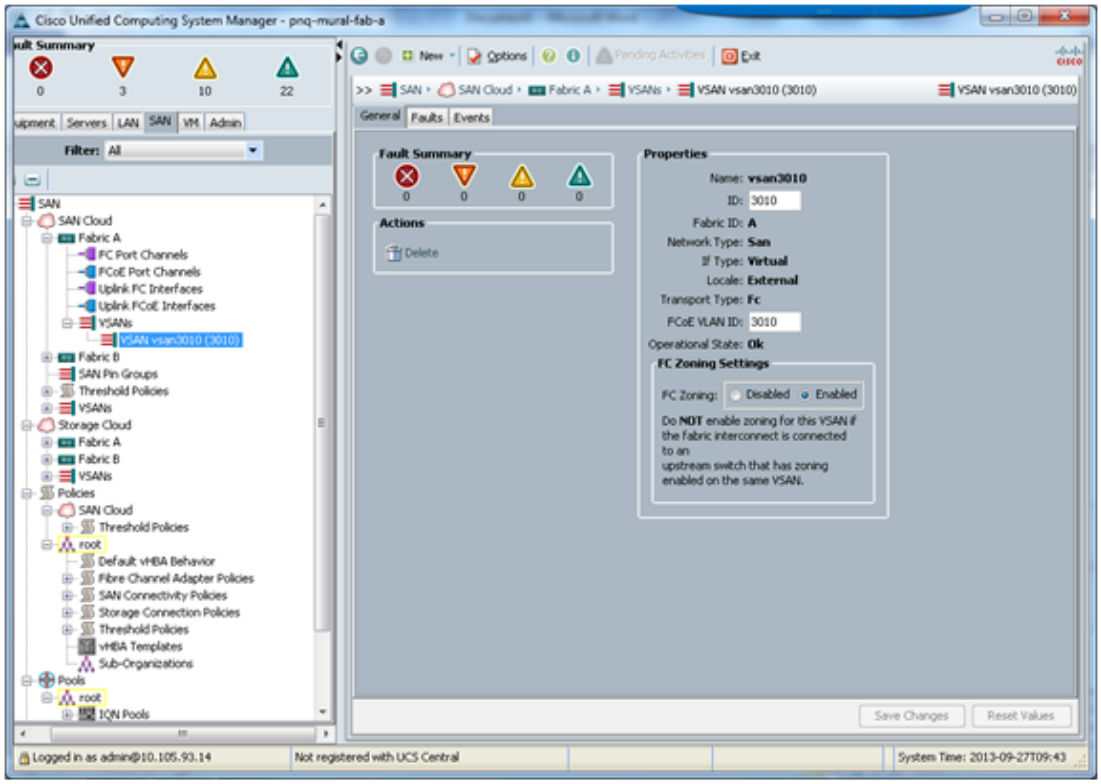
- Do not configure a VSAN as 4079. This VSAN is reserved and cannot be used in either FC switch mode or FC end-host mode.
- Do not configure VSANs with an ID in the range from 3040 to 4078 if you plan to use FC switch mode in a Cisco UCS domain. In FC switch mode, VSANs in that range are not operational and the Cisco UCS Manager marks them with an error and raises a fault.

To create a VSAN for a Fabric Interconnect:

1. In UCS Manager, navigate to the **SAN** tab > **SAN Cloud** > **Fabric A/B** > **VSANs** > right -click **VSANs** > **Create VSAN**.
2. Enter the **NAME** of the VSAN (see naming guidelines above).

3. Select **FC Zoning Enabled**, **Fabric A/B**, **VSAN ID**, and **FCoE VLAN**, as shown in the following example.

   Note: Make sure that the **FCoE VLAN ID** is same as the **VSAN ID**.



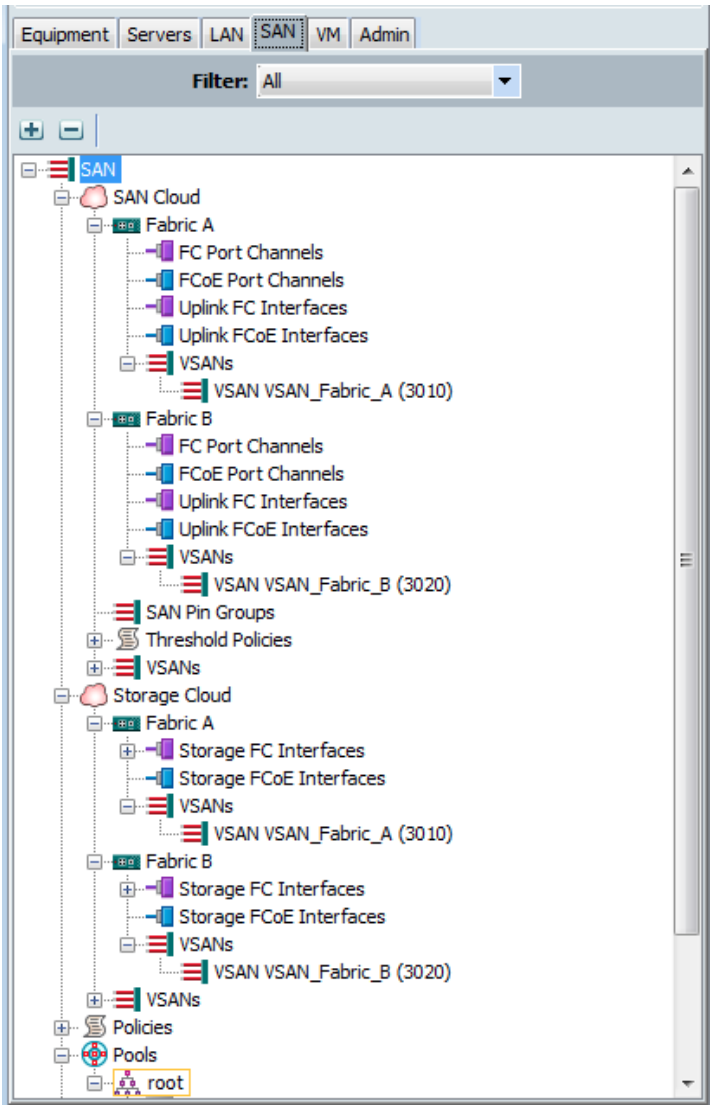
   Note: You can use the values 3010 and 3020 for the VSANs. Ensure that these values are not used in the network.

4. Repeat Step 1-3 for the other Fabric Interconnects.

5. Click **Save Changes**.

   Note: VSANs are required under both SAN and Storage clouds, as in this example for all Fabric Interconnects.

Copyright © 2014, Cisco Systems, Inc.

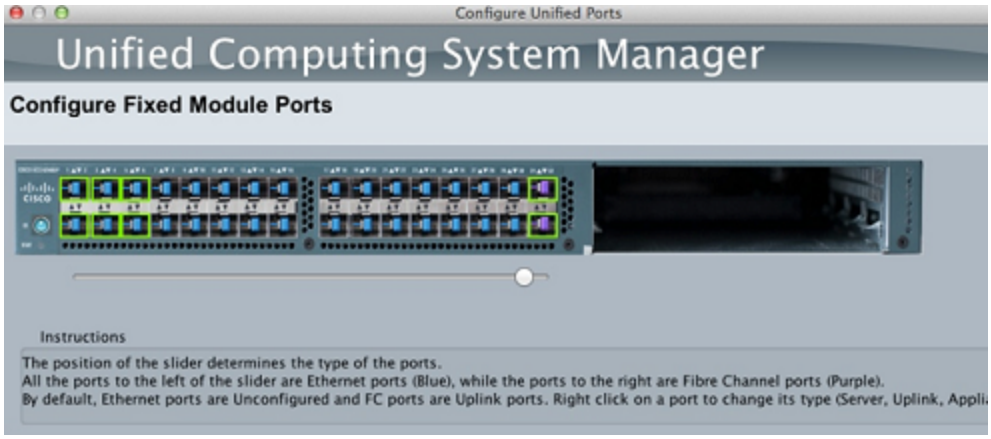


Designating Storage Ports

For each Fabric Interconnect, configure the ports connecting to the storage array as FC, reboot, and then designate the FC Ports as FC Storage Ports.

To designate storage ports:

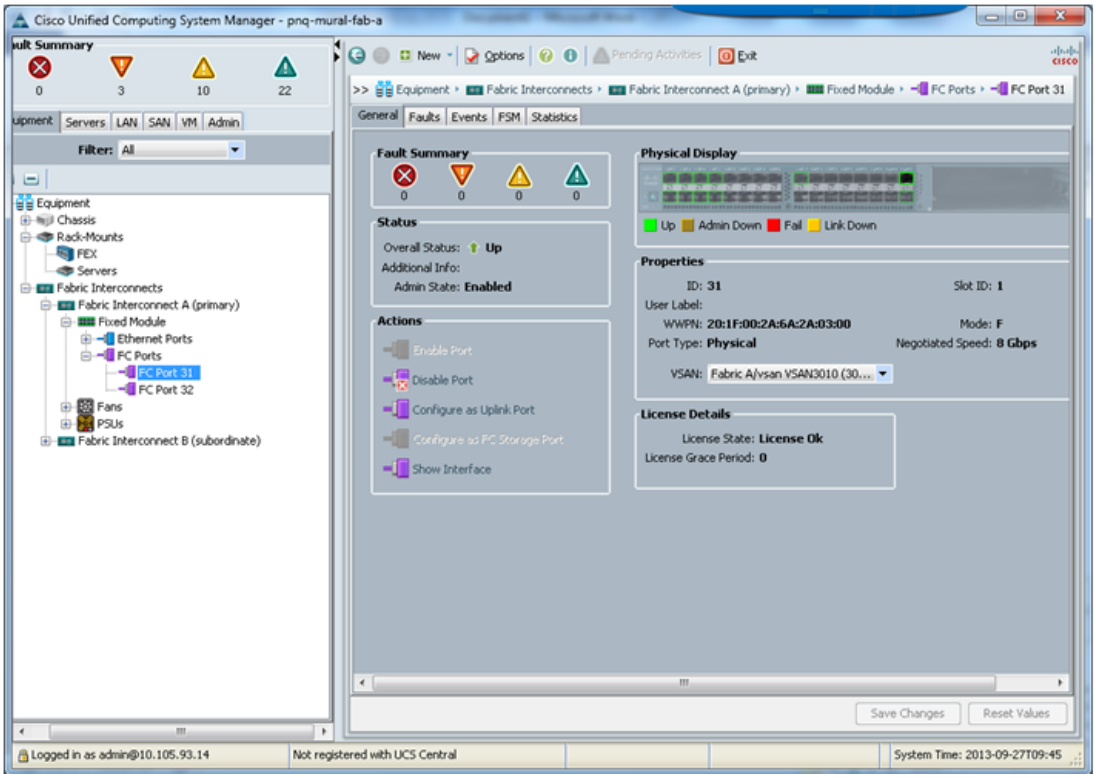
1. In UCS Manager, go to the **Equipment** tab > **Fabric Interconnect A/B** > **Configure Unified Ports**.



2. Use the slider to configure the ports connecting to the storage array as **FC**.
3. Repeat steps 1 and 2 for the other Fabric Interconnect.
4. Wait until the Fabric Interconnects have rebooted.
5. From the Equipment tab, navigate to **Fiber Interconnect** > **Fiber Interconnect A/B** > **Fixed Module** > **FC Ports** > **FC Port 31 /32**.

Copyright © 2014, Cisco Systems, Inc.

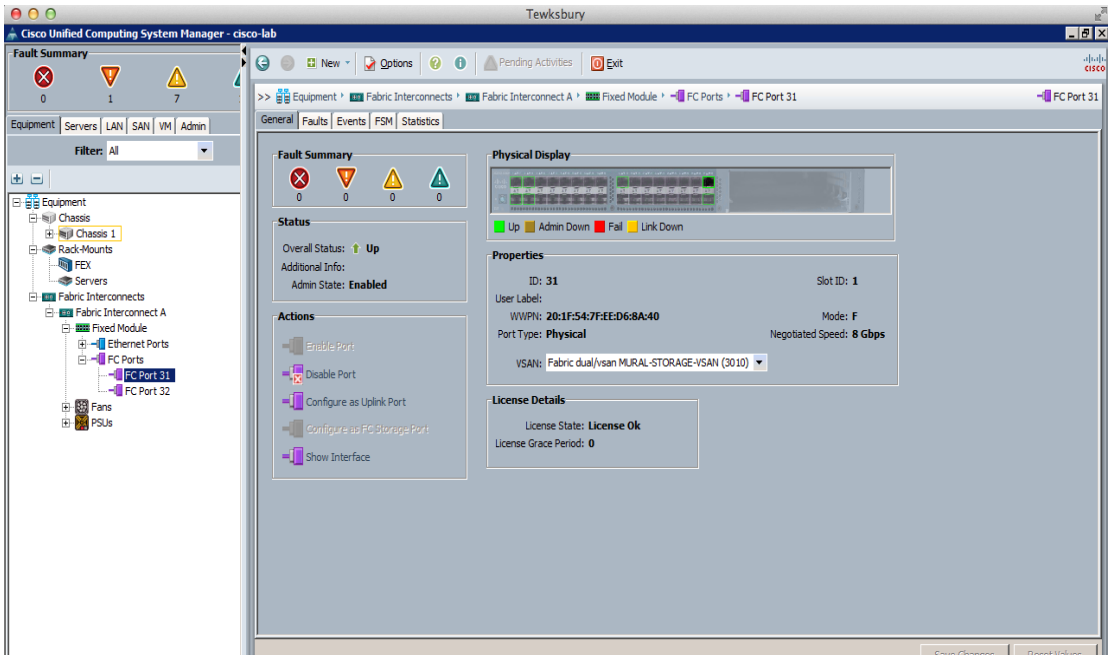
6. Select **Configure as FC Storage Port**.



Assigning the Storage Cloud VSAN to the FC Ports

To assign the Storage Cloud VSAN to the FC ports:

1. From the **Equipment** tab, navigate to **Fabric Interconnect** > **Fabric Interconnect A/B** > **Fixed Module** > **FC Ports** > **FC Port 31/32**.
2. Select the required VSAN (for example, **VSAN3010** for Fabric Interconnect A, as shown in the following image).



3. Click **Save Changes**.
4. Repeat steps 1 through 3 for the other Fabric Interconnects, ensuring that correct VSANs are selected.

Confirm Storage Port WWPN is Logged Into the Fabrics

Zoning enables access control between storage devices and user groups. Creating zones increases network security and prevents data loss or corruption. A zone set consists of one or more zones in a VSAN.

To confirm the storage port is logged into the fabrics:

1. Using SSH, log in as admin to the virtual IP address for the UCS Fabric Interconnect.
2. Run the **connect** command to enter the NX-OS CLI:

   cisco-lab-A# **connect nxos**

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2013, Cisco Systems, Inc. All rights
```


Copyright © 2014, Cisco Systems, Inc.

```
reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed
under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of
each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

cisco-lab-A(nxos)#
```

Note: You can use the **?** option at the prompt to get help.

3. View the active zonesets:

   cisco-lab-A(nxos)# **show zoneset active**

```
zoneset name ucs-cisco-lab-vsan-3010-zoneset vsan 3010
zone name ucs_cisco-lab_A_8_UI1_vHBA-A vsan 3010
* fcid 0x6c0000 [pwwn 20:00:00:05:ad:1e:11:2f]
* fcid 0x6c00ef [pwwn 50:06:01:68:3e:e0:0a:6b]
* fcid 0x6c01ef [pwwn 50:06:01:60:3e:e0:0a:6b]

zone name ucs_cisco-lab_A_7_GMS1_vHBA-A vsan 3010
pwwn 20:00:00:05:ad:1e:11:4f
* fcid 0x6c00ef [pwwn 50:06:01:68:3e:e0:0a:6b]
* fcid 0x6c01ef [pwwn 50:06:01:60:3e:e0:0a:6b]

zone name ucs_cisco-lab_A_6_INSTA1_vHBA-A vsan 3010
* fcid 0x6c0004 [pwwn 20:00:00:05:ad:1e:11:7f]
* fcid 0x6c00ef [pwwn 50:06:01:68:3e:e0:0a:6b]
* fcid 0x6c01ef [pwwn 50:06:01:60:3e:e0:0a:6b]
```

```
zone name ucs_cisco-lab_A_5_INSTA2_vHBA-A vsan 3010
* fcid 0x6c0005 [pwwn 20:00:00:05:ad:1e:11:5f]
* fcid 0x6c00ef [pwwn 50:06:01:68:3e:e0:0a:6b]
* fcid 0x6c01ef [pwwn 50:06:01:60:3e:e0:0a:6b]
```

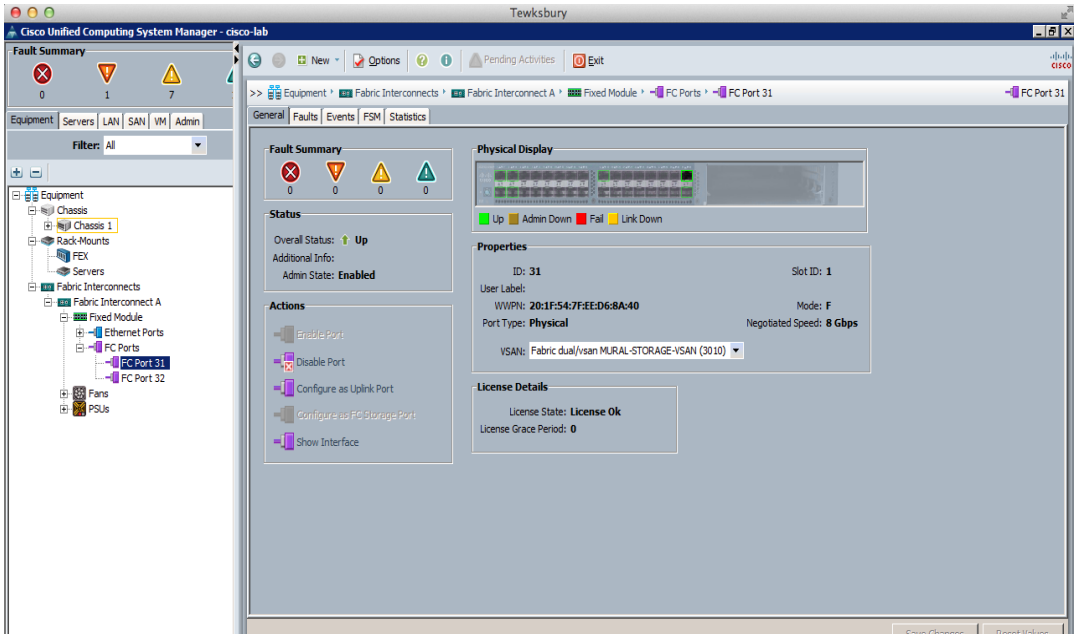
4. Run the `show flogi database vsan` *`vsan-ID`* command, where *vsan-ID* is the identifier for the VSAN. In this example, 3010 is used for VSAN Fabric Interconnect A.
5. Make a note of the WWPNs as in this example. These WWPNs will be required when creating storage connection policies:

```
mural-norse-lab-B (nxos)# show flogi database vsan 3010
INTERFACE     VSAN      FCID        PORT NAME                NODE NAME
--------------------------------------------------------------------------
fc1/31        3010      0x1000ef    50:06:01:60:3e:a0:28:d2
50:06:01:60:be:a0:28:d2
fc1/32        3010      0x1001ef    50:06:01:69:3e:a0:28:d2
50:06:01:60:be:a0:28:d2
```

6. Run the `exit` command.
7. Enter the connect **nxos B**, where B is for Fabric Interconnect B.
8. Enter `show flogi database vsan` *`vsan-ID`*, where *vsan-ID* is the identifier for the VSAN. In this example, we use 3020 for VSAN on Fabric Interconnect B.
9. Make a note of the WWPN.

```
mural-norse-lab-B(nxos)# show flogi database vsan 3020
--------------------------------------------------------------------------
INTERFACE    VSAN    FCID        PORT NAME                  NODE NAME
fc1/31       3020    0x4200ef    50:06:01:61:3e:a0:28:d2
50:06:01:60:be:a0:28:d2
fc1/32       3020    0x4201ef    50:06:01:68:3e:a0:28:d2
50:06:01:60:be:a0:28:d2
```

 Copyright © 2014, Cisco Systems, Inc.



Creating Storage Connection Policies

Create a storage connection policy for each Fabric Interconnect.

To create storage connection policies:

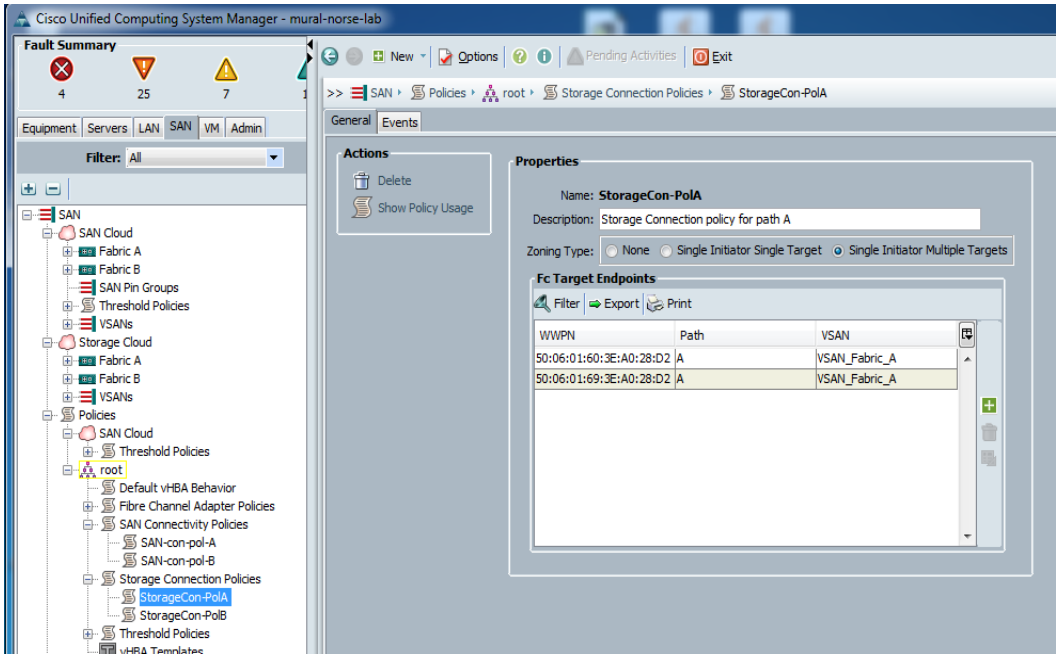
1. From the **SAN** tab, navigate to **Policies** > **root**, right-click on **Storage Connection Policies**, and select **Create Storage Connection Policies**.
2. Enter a name that complies with local naming conventions and the values as shown below:

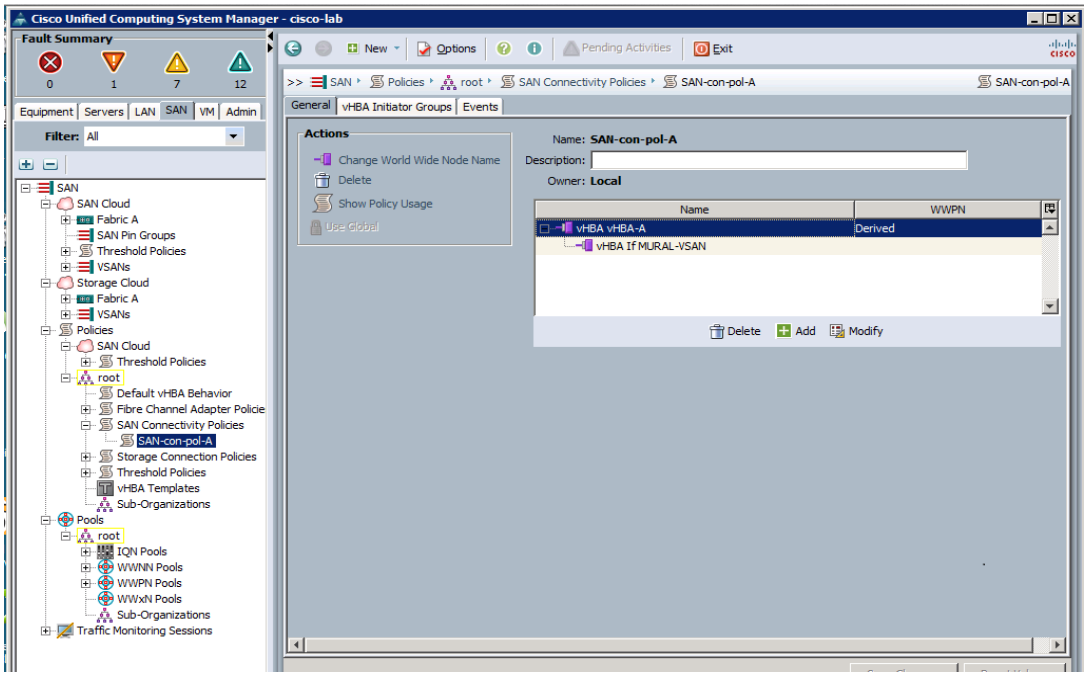
```
Name: storage-conn-polA
Zoning = Single Initiator Multiple Target

Add FC Target Endpoints
In FC Target Endpoints, enter WWPN printed in the last section
for the related VSAN.
Path (select FI A or B)
VSAN (created earlier for respective FI)
```

3. Repeat Step 2 for the other WWPN. This will create another FC Target Point.

4. Repeat Steps 1 through 4 to create storage connection policies for the other Fabric Interconnect. The following example screen shots show two Fabric Interconnects, FI-A and FI-B, each with two Target Endpoints.





Copyright © 2014, Cisco Systems, Inc.

Creating SAN Connectivity Policy

A virtual host bus adapter (vHBA) logically connects a virtual machine to a virtual interface on the UCS 6100 series Fabric Interconnect and allows the virtual machine to send and receive traffic through that interface.You must create a vHBA initiator group for each vHBA.

Connectivity policies determine the connections and the network communication resources between the server and the LAN or SAN on the network. These policies use pools to assign MAC addresses, WWNs, and WWPNs to servers and to identify the vNICs and vHBAs that the servers use to communicate with the network.

If you want to support any VSAN, it needs to be configured globally into Cisco UCS Manager, and then it can be associated with a particular vHBA.

To create a vHBA initiator group for the storage connectivity policy:

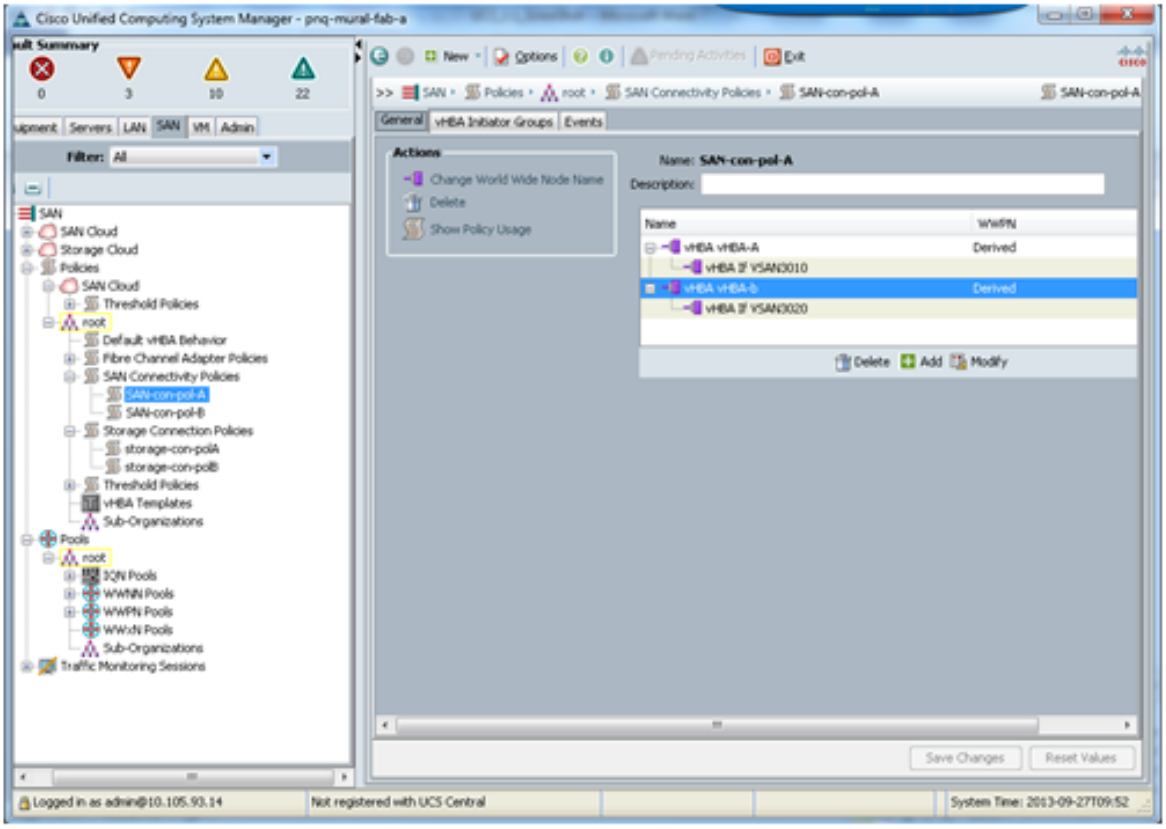
1. From the **SAN** tab, navigate to **Policies** > **root**.
2. Right-click **SAN Connectivity Policies** and select **Create SAN Connectivity Policy**.
3. Click **Add**. Enter the values shown in the following example, ensuring that the name complies with local naming conventions.

```
Name : vHBA-A
WWNN Assignment: wwnn-pool1(768/784)
Fabric ID : A
Select VSAN: VSAN3010 (Storage Connection Policy (created
earlier)
Adaptor Policy : VMWare
```

4. Repeat Steps 1 through 3 for the other vHBA.

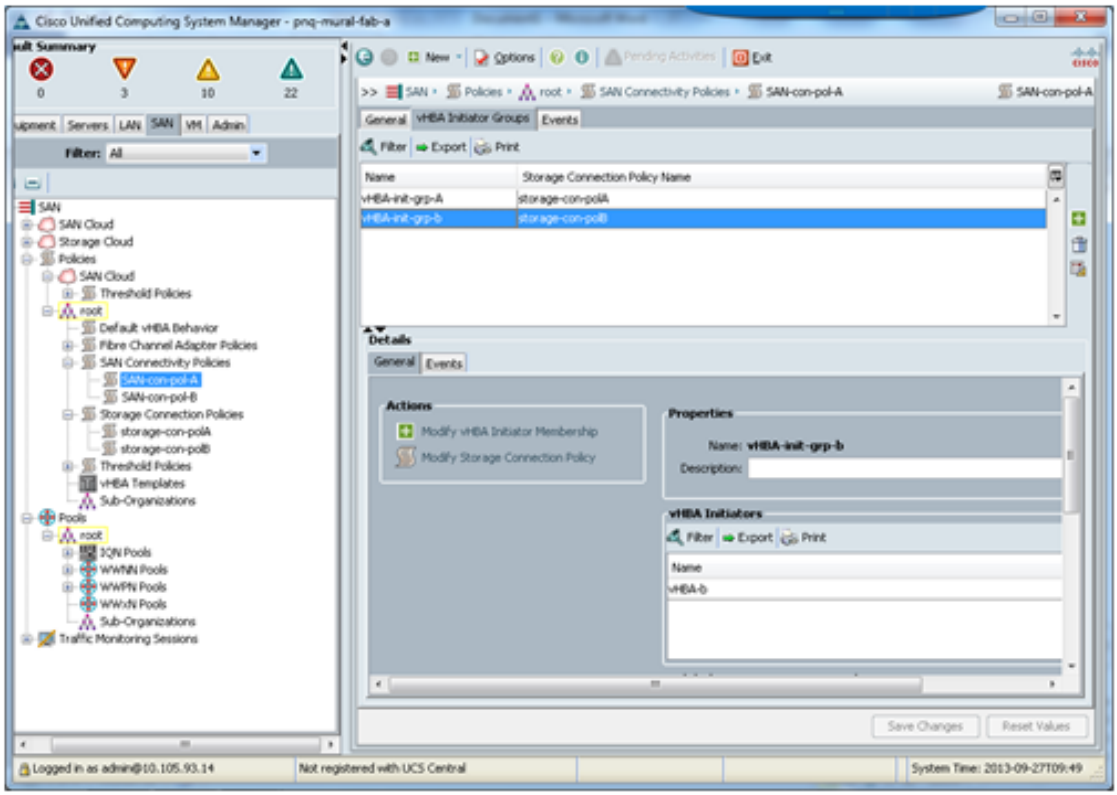
Configuring SAN Cloud Policy

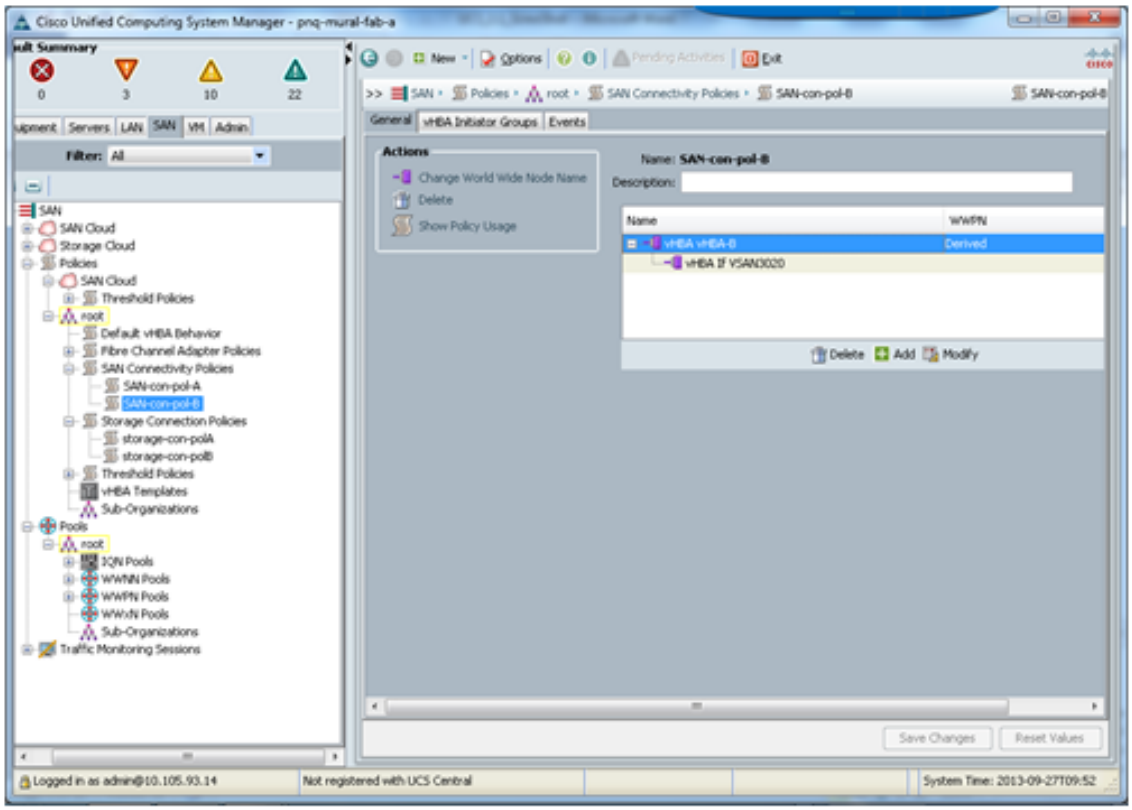
The SAN cloud policy **San Con Pol A** shows two vHBAs: **vHBA-A** and **vHBA-b**:

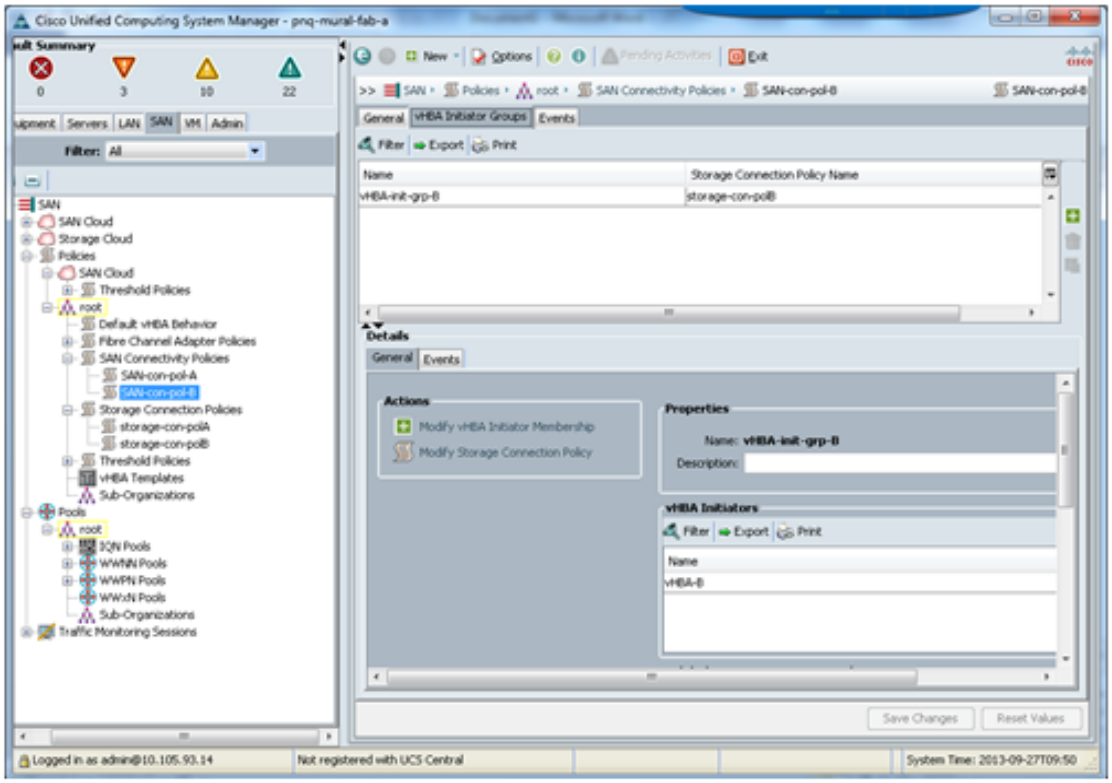


The following screen shows an example of two vHBAs initiator groups within one SAN connection policy. The section after the screen details the steps for creating a VHBA initiator group.

Copyright © 2014, Cisco Systems, Inc.


Cisco Unified Computing System Manager - pnq-mural-fab-a
SAN
SAN Cloud
Storage Cloud
Policies
Threshold Policies
root
Default vHBA Behavior
Fibre Channel Adapter Policies
SAN Connectivity Policies
SAN-con-pol-A
SAN-con-pol-B
Storage Connection Policies
storage-con-polA
storage-con-polB
vHBA Templates
Sub-Organizations
Pools
IQN Pools
WWNN Pools
WWPN Pools
WWxN Pools
Traffic Monitoring Sessions
vHBA-init-grp-A
vHBA-init-grp-b
Modify vHBA Initiator Membership
Modify Storage Connection Policy
Name: vHBA-init-grp-b
vHBA-b
Save Changes
Reset Values
Logged in as admin@10.105.93.14
Not registered with UCS Central


Cisco Unified Computing System Manager - pnq-mural-fab-a
SAN
SAN Cloud
Storage Cloud
Policies
Threshold Policies
root
Default vHBA Behavior
Fibre Channel Adapter Policies
SAN Connectivity Policies
SAN-con-pol-A
SAN-con-pol-B
Storage Connection Policies
storage-con-polA
storage-con-polB
vHBA Templates
Sub-Organizations
Pools
IQN Pools
WWNN Pools
WWPN Pools
WWxN Pools
Traffic Monitoring Sessions
Change World Wide Node Name
Delete
Show Policy Usage
Name: SAN-con-pol-B
vHBA vHBA-B
vHBA IF VSAN3020
Derived
Add
Modify
Save Changes
Reset Values
Logged in as admin@10.105.93.14
Not registered with UCS Central



Create vHBA Initiator Groups

To create a vHBA initiator group for the storage connectivity policy:

1. From the **SAN** tab, navigate to **Policies** > **Root** > **SAN Connectivity Policies**.
2. Add SAN Connectivity Policies for FI A/B (above example shows **SAN-con-pol-A**).
3. Select **SAN Connectivity Policies** (Example: SAN-con-pol-A).
4. Add values as shown below:

```
Name: vHBA-init-grp-A
Select vHBA Initiators (for example, vHBA-B)
Storage Connection Policy (for example, Storage-con-polB)
```

5. Click **OK** to save changes.
6. Repeat the above steps for the other Fabric Interconnect.

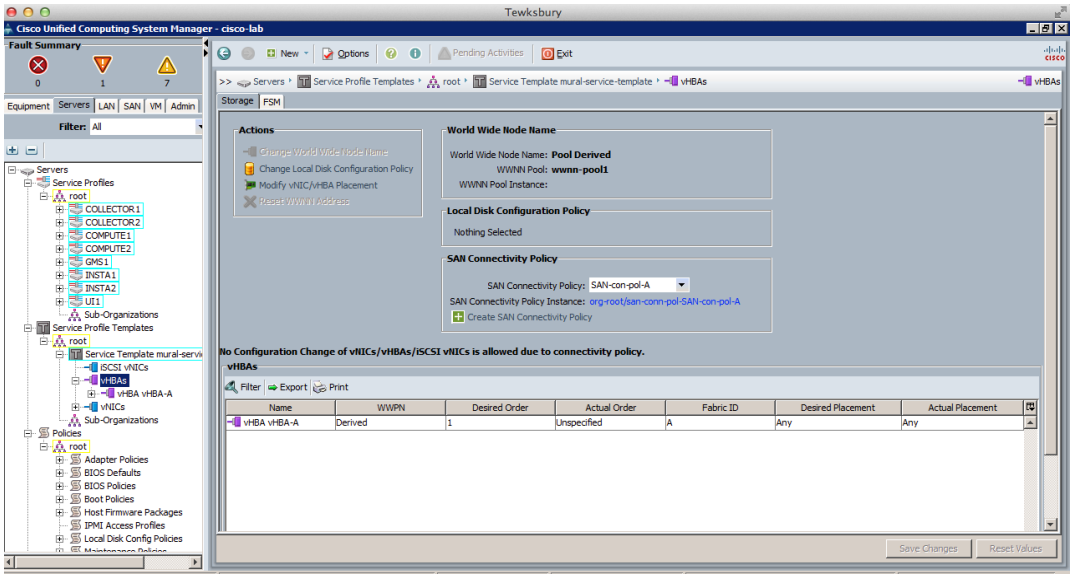
Copyright © 2014, Cisco Systems, Inc.

Verify Service Profile Templates

When vHBA initiator groups are created, vHBAs are updated into service profile templates.

To verify service profile templates:

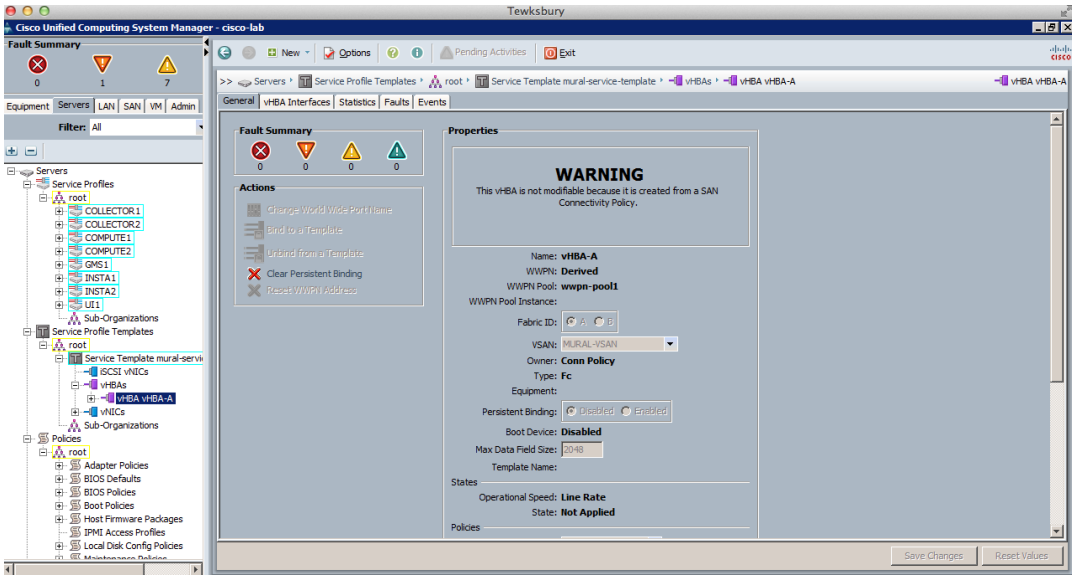
1. Navigate from the **Servers** tab to **Service Profile Templates** > **root** > **Service Template** (for example, Mural-service-template) >**VHB**.
2. Select **SAN Connectivity Policy**



3. Verify that vHBAs have been applied to a service profile template, and that all details are correct.
4. Click **Save Changes**.
5. Repeat steps 1 through 4 for the second vHBA.

   The following screen shot is an example of a vHBA configuration within a

service template.



This completes the configuration of Cisco UCS Direct-Attached SAN.

System Profile Settings for the Cisco UCS Manager

You can configure profiles in the Cisco UCS Manager for modifying the default hardware settings of the systems. The following are recommended settings for the MURAL platform.

Ethernet Adapter Policy

The Ethernet adapter policy for all Ethernet interfaces requires the values listed in the following table. Fields that are not mentioned, or whose values are specified as *Default*, need not be changed from the default values.

To configure Ethernet adapter policy on the UCS:

1. Navigate to **Servers** > **Policies** > **root** > **Eth Adapter Policy Default** > **General**.
2. Enter the field values as shown in the following table.

   Note: Ensure that the Resources and Options values are set correctly, as recommended in the table.

Copyright © 2014, Cisco Systems, Inc.

Field Name	Field Value
Transmit Queues	1
Ring Size	*Default*
Receive Queues	8
Ring Size	*Default*
Completion Queues	9
Interrupts	16
Transmit Checksum Offload	Enabled
Receive Checksum Offload	Enabled
TCP Segmentation Offload	Enabled
TCP Large Receive Offload	Enabled
Receive Side Scaling	Enabled
Fallback Timeout	*Default*
Interrupt Mode	MSI X
Interrupt Coalescing type	Min
Interrupt Timer	350

Bios Policy (Processor Page)

The **Processor** page under **Bios Policy** requires the values listed in the following table. Fields that are not mentioned, or whose values are specified as *Default*, need not be changed from the default values.

To configure Bios Policy on the UCS, navigate to **Servers** > **Policies** > **root** > **BIOS Policies** > **mural-bios** > **Advanced** > **Processor** and enter the field values, as shown in the following table:

Field Name	Field Value
Turbo Boost	Disabled
Enhanced Intel SpeedStep	Enabled

Field Name	Field Value
Hyper Threading	*Default*
Core Multiprocessing	all
Execute Disable Bit	*Default*
Virtualization Technology	Disabled (Enabled if VMs are expected to be run on the systems)
Direct Cache Access	Enabled
Processor C State	Disabled
Processor C1E	Enabled
Processor C3 Report	*Default*
Processor C6 Report	*Default*
Processor C7 Report	*Default*
CPU Performance	hpc
Max variable MTRR Setting	*Default*

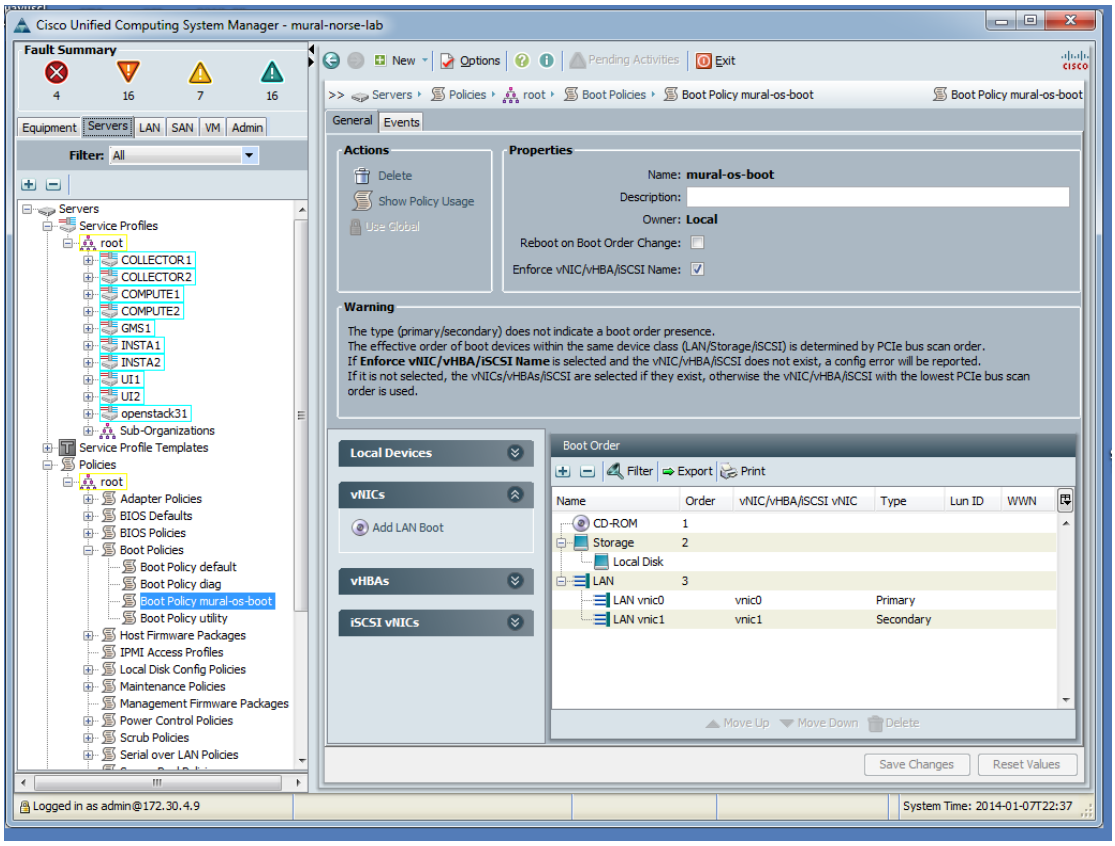
Specifying Boot Order of Devices

To specify boot order for devices, navigate to **Servers** > **Policies** > **root** > **Boot Policies** > **Boot Policy Default (or configured for service profiles)** > **General**.

Set the order of preference for boot devices as follows, as in the sample screen below:

1. Local Disk
2. LAN Boot (in the case of PXE boot, ensure that both interfaces are added--for example, **vnic0** and **vnic1**)
3. Other Devices

Copyright © 2014, Cisco Systems, Inc.



Setting the Raid Policy

We recommend that you run a RAID 1 setup. Cisco systems allow configuration of RAID setup using the LSI MegaRaid Configuration tool in BIOS.

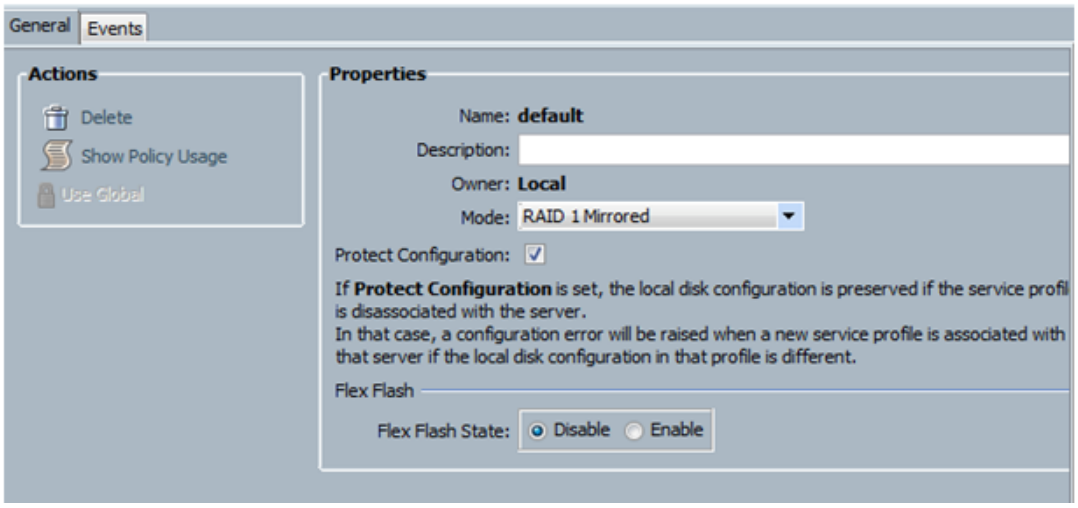
Note: Set up hardware RAID 1 on the local disk.

Caution: Do not use the **Any** configuration mode on servers with MegaRAID storage controllers.

To set the RAID policy:

1. Navigate to **Servers** > **Policies** > **root** > **Local Disk Configuration Policies** > **Local Disk**.
2. Select **Configuration Policy Default** > **General**.

3. Select **Mode: RAID 1 Mirrored**, as shown below:



For more information, see the *Cisco UCS Manager GUI Configuration Guide, Release 2.1*.

Copyright © 2014, Cisco Systems, Inc.

Manufacturing the GMS Blade

The GMS blade hosts the GMS node, a MURAL component that enables centralized installation and monitoring of all the blades on the MURAL system.

Follow these steps to manufacture (install the operating system software on) the GMS blade.

Before you begin:

- Configure Serial over LAN (SOL) on all the blades during EMC setup.
- Locate your CIQ, and refer to it for such details as UCS access credentials and KVM SOL IP address.

To manufacture the GMS blade:

1. Download the ISO image that was included with the MURAL software package to the machine from which you will access the Cisco UCS blades.

   The ISO image file is:
   `md5sum /data/mfgcd-guavus-x86_64-20130925-174106.iso`
   The MD5 checksum for the image is:
   `48fc39f77bc7ab8847ca2f7a684d2ace /data/mfgcd-guavus-x86_64-20130925-174106.iso`

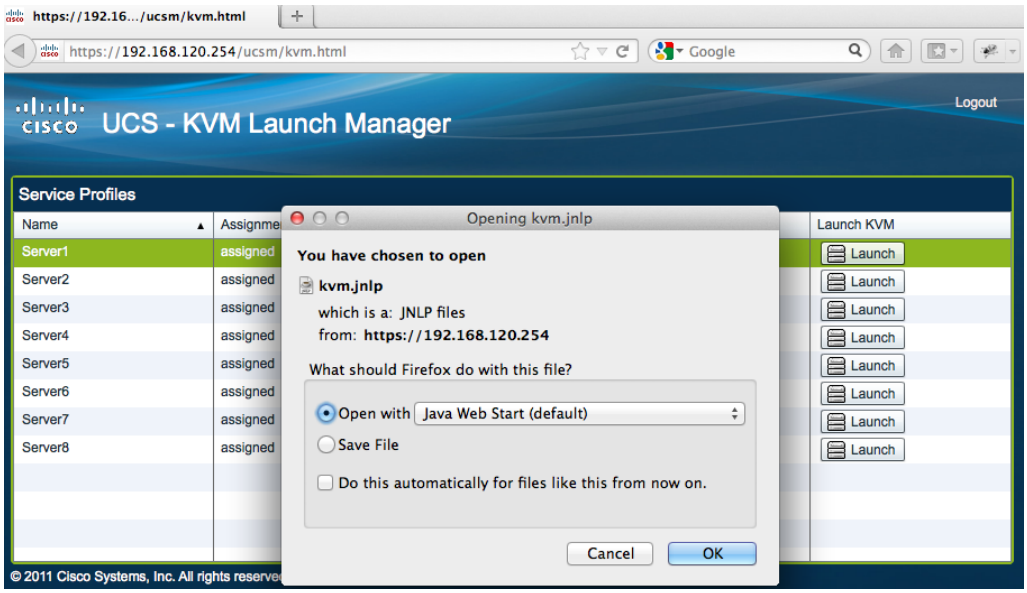
2. Open the Cisco UCS - KVM Launch Manager in a browser and enter your login credentials.

   Note: Firefox with Java version 6, or above, is the most suitable browser to access the UCS.

   You will see all the blades available on the chassis.

3. Click the **Launch** button for the first node (Server1).

   Click **OK** to download and open a **kvm.jnlp** file.



A keyboard access warning message appears—click **OK**.

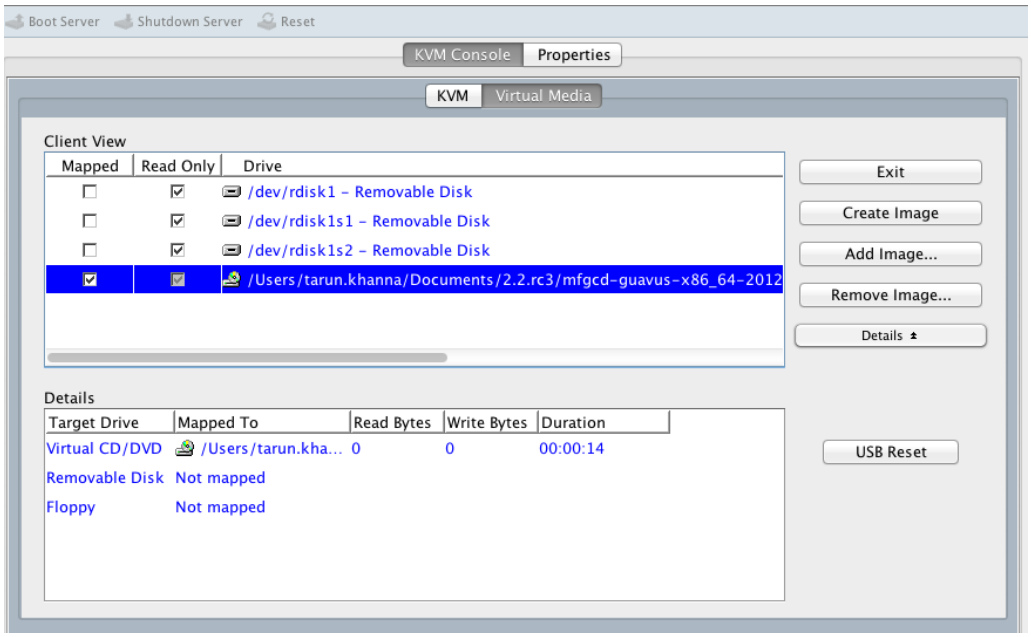
The console for the port opens.

4. From the **Virtual Media** tab, click **Add Image**, and specify the path of the ISO image that you downloaded in Step 1.

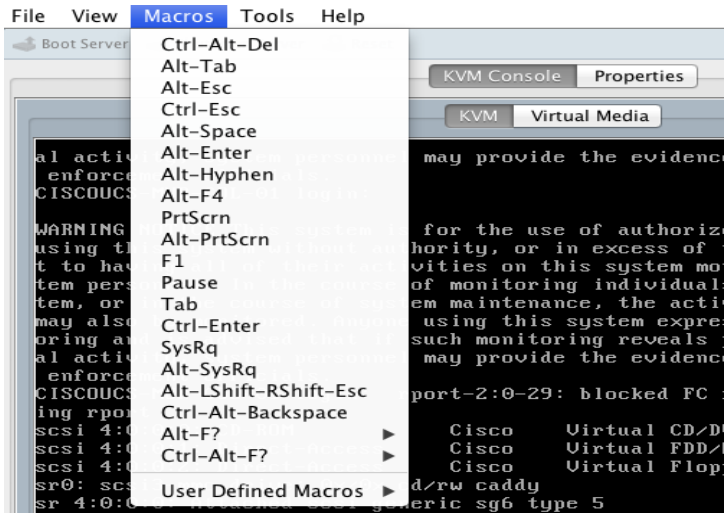
   Note: Copy the ISO image to multiple directories if you need to manufacture more than one blade at the same time. Select the ISO image from these copied directories for each individual node. Be sure that the same ISO image is not selected for two nodes while manufacturing is in progress for a given node.

5. Mount the ISO image by clicking the check box next to the added image.

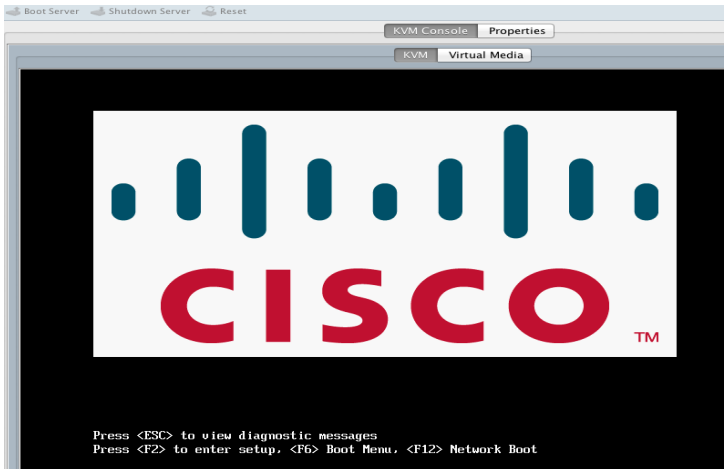
Copyright © 2014, Cisco Systems, Inc.



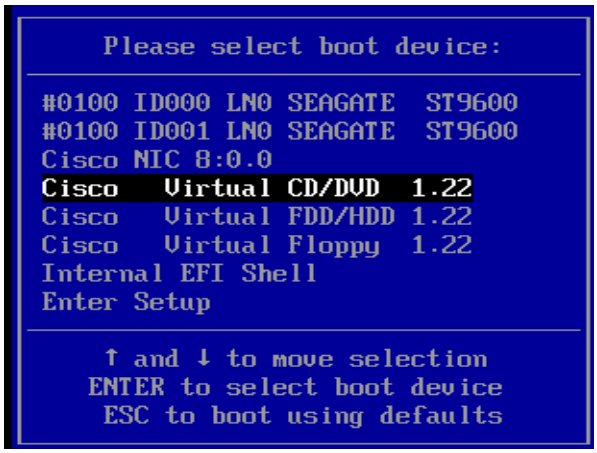
6. Reboot the blade to boot with the newly mounted image. Click the **KVM** tab and select **Ctrl-Alt-Del** from the **Macros** drop-down menu by hovering over the option.



7. When the boot order screen appears, press **F6** to select the boot order.



8. Select **Virtual CD/DVD** so the blade boots with the mounted ISO image.



9. At the # prompt, start the manufacture process by running the `manufacture` command:

   For 1 Disk : In case of H/W Raid 1

   ```
   # manufacture.sh -v -t -f /mnt/cdrom/image.img –m 1D –L 1D --cc no
   --cs no --cl no –a
   ```

   Note: In the above command variable, `-L 1D` specifies the use of one hard disk in a blade. Please ensure that two hard disks are available in the blade, one for each side of the mirrored RAID setup.

```
Running  /etc/in it .d/rcS .d/S34automf g
– Automatic manufacture is not enabled. Type ‘automf g’ to start it.
```


Copyright © 2014, Cisco Systems, Inc.

```
BusyBox v1.00 (2010.12.03—23:16+0000) Built—in shell (ash)
Enter ‘help’ for a list of built—in commands.

—sh: can’t access tty: job control turned off

Processing  /etc/profile... Done

# manufacture.sh —v —t —f /mnt/cdrom/image.img
======  Starting manufacture at 20120730—091426
======  Called as: /sbin/manufacture.sh —v —t —f /mnt/cdrom/image.img
================================================
Manufacture script starting
================================================
---------------------------------
Model selection
 ---------------------------------

Product model (‘?‘ for info) ( DESKTOP VM VM_2D 1D 2D 3D 4D 2D_EXT4)
[2D]: 2D

== Using model: 2D

---------------------------------
Kernel type selection
---------------------------------
Kernel type (uni smp) [smp]:
== Using kernel type: smp
---------------------------------
Layout selection
---------------------------------
Layout (STD) [2D]:
== Using layout: 2D
---------------------------------
Partition name—size list selection
---------------------------------

Partition name—size list [VAR 40960 SWAP 40960]:
```

```
== Using partition name—size list: VAR 40960 SWAP 40960


---------------------------------
Device list selection
---------------------------------
Device list [/dev/sda   /dev/sdb]
== Using device list: /dev/sda   /dev/sdb


---------------------------------
Interface list selection
---------------------------------
Interface list [eth0  eth1]:
== Using interface list: eth0  eth1


---------------------------------
Interface naming selection
---------------------------------
Interface naming [none]:
== Using interface naming: none
== Smartd enabled
---------------------------------
CMC server settings
---------------------------------
Enable CMC server (yes no) [yes]: no
Enable CMC server (ves no) [ves]: no
== CMC server enabled: no


---------------------------------
CMC client settings
---------------------------------
Enable CMC client (ves no) [ves]: no
== CMC client enabled: no
== CMC client auto—rendezvous enabled: no
== CMC server address for rendezvous: (none)


---------------------------------
Cluster settings
```

 Copyright © 2014, Cisco Systems, Inc.

```
---------------------------------
Enable cluster (yes no) [no]: no
== Cluster enable: no
Cluster ID:

Number Start End Size File system Mame Flags
1 O.OZMiB 30000MiB 30000MiB ext3 primary
Writing partition table to DISICZ
Disk devsdb: S7Z3Z6MiB
Sector size (logicalphysical): 512B,51ZB
Partition Table: gpt
Number Start End Size File system Mame Flags
1 O.OZMiB 5?Z3Z6MiB 5?ZJZ6MiB ext3 primary

==== Making filesystems
== Creating  ext3  filesystem on  /dev/sda2 for  BOOT1
== Creating  ext3  filesystem on /dev/sda3  for BOOT2
== Creating  ext3  filesystem on /dev/sda1 for BOOTMGR
== Creating  ext3  filesystem on /dev/sda8  for CONFIG
== Nothing  to do on /dev/sda9  for  HA
== Creating  ext3  filesystem on /dev/sda5  for ROOT1
== Creating  ext3  filesystem on /dev/sda6  for ROOT2
== Making swap on /dev/sda7  for SWAP
Setting up swapspace version 1, size = 42952404 kB
== Creating  ext3  filesystem on /dev/sda10  for VAR

== CMC server address for rendezvous: (none)

---------------------------------
Cluster settings
---------------------------------
== Cluster  enable: no
Cluster ID:
== Cluster ID: (none)
Cluster  description:
== Cluster  description: (none)
Cluster interface:
```

```
== Cluster  interface: (none)
Cluster master virtual IP address [0.0.0.0]:
== Cluster  master  virtual  IP address: 0.0.0.0
Cluster master virtual IP masklen [0]:
== Cluster master virtual IP masklen:  0
Cluster shared secret :
```

The system shows a message **Manufacture done** and returns to a **#** prompt when the manufacturing of a blade is complete.

10. Deselect the ISO image selected in Step 5 of this procedure. Type **reboot** to reboot the node with the new ISO image.

```
== System successfully imaged
—— Writing Host ID: 3b0455ef813d
== Zeroing the destination partition disk /dev/sda9 with dd
== Calling imgverify to verify manufactured system
== Using layout: ZD
Using dev list:  /dev/sda /dev/sdb
== Verifying image location 1
=== Mounting partitions
=== Checking manifest
=== Unmounting  partitions
=== Image location 1 verified successfully.
== Verifying image location 2
=== Mounting partitions
=== Checking manifest
=== Unmounting partitions
=== Image location 2 verified successfully.
== Done
======  Ending manufacture at 20120730—105002
—— Manufacture done.
#
#
# reboot
```

11. Log in to the GMS blade (using SSH) with the username **admin**. The system prompt will be presented. Continue to the next section, "Setting Up the GMS Node" on the facing page.

 Copyright © 2014, Cisco Systems, Inc.

Setting Up the GMS Node

After the GMS node is manufactured (the MURAL software is loaded) set up the GMS node by setting the password for the node, assigning IP addresses, downloading any available patches, and starting the server, as described in the following procedure.

To set up the GMS node:

1. Log in to the GMS node using the console.
2. Set the password:

```
> en
# conf t
(config) # username admin password admin@123
(config) # write mem
(config) # exit
```

3. Assign the IP addresses for the management interface and default gateway:

```
> en
# conf t
(config) #interface mgmt-interface ip-address mgmt-IP-of-GMS-
server  subnetmask-of-mgmt-net
(config) # ip-default-gateway mgmt-network-default-gateway-IP
(config) # license install LK2-RESTRICTED_CMDS-88A4-FNLG-XCAU-U
(config) # write memory
(config) # _shell
```

4. Download the patches from the FTP server to the GMS server in the **/data** directory. Apply all patches applicable for GMS nodes.

   Note: See the release notes for software version 3.2 MR2 (or later), for a complete list of patches and installation instructions.

5. Start the GMS server:

```
> en
# conf t
(config) # pm process gms_server launch auto
(config) # pmx register pgsql
```

```
(config) # write memory
(config) # pm process tps restart
(config) # pm process gms_server restart
```

6. Check the status of the GMS server by running the following command:

```
> en
# _shell
# cli -t "en" "config t" "show pm process gms_server" | grep
"Current status"
```

7. Proceed to "Defining the Deployment Topology Using GMS" on the facing page

 Copyright © 2014, Cisco Systems, Inc.

Defining the Deployment Topology Using GMS

This section describes how to define the deployment topology using the General Management System (GMS). GMS is a node that enables you to do a centralized installation, rather than manually installing and configuring software on each blade.

Note: The GMS interface requires Java Plug-in 10.40.2.43 or above, and JRE version 1.7.0_40-b43 or above.

When you first configure the GMS, WWIDs for the LUNs are not yet configured (because you have not yet configured EMC) and you must use dummy values as placeholders. You will configure the LUNs in "Setting Up Data Storage On the EMC" on page 80, and then revisit this section to specify the actual WWIDs.

Note: The **VLANs** and **Bonds** tabs on the GMS interface are not used in this version of the MURAL application.

Before you begin, locate your Customer Information Questionaire (CIQ)—for some of the steps in this procedure, you will need to refer to the completed information from the **IP Survey** tab.

Accessing the GMS Interface

To access the GMS interface:

1. Enter the following URL from a browser window:

   http://*IP-Address-GMS*/configure

2. Enter the login name and password that you configured previously (see "Setting Up the GMS Node" on page 53).

3. In the Java applet pop-up window, enter your user name and password:



4. Proceed with the initial configuration of the system as outlined in the following sections.

Loading the Base Configuration File

A sample configuration file, called **mural.xml**, for the MURAL software version is provided with the MURAL installation package. It includes a logical grouping of the nodes into two chassis (blade enclosure) and configuration settings for all of the system components based on the details of your network that were supplied in the CIQ for your environment.

Note: We strongly recommend that you use the GMS interface for all modifications to the **mural.xml** file, and that you do not modify the file directly.

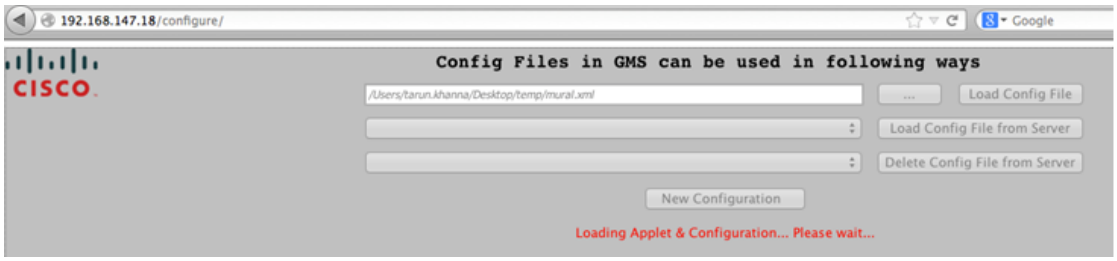
The settings for the MURAL system are preconfigured based on the **mural.xml** file, and in most cases you only need to verify these settings. There are cases where you will need to modify the settings, for example, if the slot numbers for some of the blades was changed or if there were values that were not yet known when you filled out the CIQ. Using the GMS interface, you can modify the settings for the system elements.

Profiles are provided for system components, and in some cases you might select from among different profiles to choose the one that provides optimal settings for your environment.

To upload the configuration file from the local machine to the GMS:

1. From the GMS interface, click the file navigation button (**...**) and navigate to the configuration file, then click **Load Config File**.

 Copyright © 2014, Cisco Systems, Inc.



2. Upon completion of the load process, the GMS interface appears. The fields within each tab are populated according to the information from the **mural.xml** configuration file, such as hostnames and slot numbers for each of the nodes.

   Because the fields are pre-populated based on a configuration file that is specific to your deployment, for many settings you need only verify that the setting is correct.

   Proceed with the initial configuration as outlined in the following sections.

Verifying Server Details for Each Node

From the **Server Details** tab, verify the MAC addresses of interface bindings and the WWIDs for the storage that is to be allocated to each node—master and backup Collector nodes, all Compute nodes, both UI/Caching (Rubix) nodes, both Insta nodes, as well as the MAC address of the GMS node. During the initial configuration process, you must do this for each node in each of the two logical chassis.

Finding MAC Addresses for Nodes

To find the MAC address for each node (blade), login to the UCS Manager and go to **Server** > **root**, and select and expand the blade. For example, click on `COLLECTOR1/Expand Vnics/Vnic0` or `Vnic1` to find the MAC addresses for interfaces **eth0** and **eth1**, respectively.

Finding WWIDs of LUNs

If you are at a stage where UCS configuration is not yet fully set up, then you can use a dummy WWID for each blade. However, if you have configured both UCS

and EMC, you must enter the actual WWID, which you can find as described below.

To find the WWID of LUNs:

1. Login to the EMC interface at **http://*IP-Address-EMC***.
2. From the **Storage** tab, click **LUNs**, highlight the destination LUN, and click **Properties**.
3. The **unique ID** is displayed. To obtain the WWID, remove the separator '**:** ' from the unique ID and prefix the complete unique ID with 3, for example:

   ```
   360060160c7102f004887f4015d49e211
   ```

In MURAL, the number of LUNs, and therefore the number of WWIDs, assigned to each node depends on the type of node:

- Collector nodes, Compute nodes, and UI nodes each have one WWID assigned.
- Each Insta node (being clustered) has three assigned WWIDs—one assigned to **dbroot1** of both Insta 1 and Insta 2, one assigned to **dbroot2** of both Insta 1 and Insta 2, and one **pgsql**. The **pgsql** storage partitions are for the Postgres database and will store Rubix-related data, and there is one pgsql storage partition for each Insta node.

Important: Ensure that the same WWID is assigned to **dbroot1** for both Insta 1 and Insta 2 in the GMS configuration. Likewise, ensure that the same WWID is assigned to **dbroot2** for both Insta nodes.

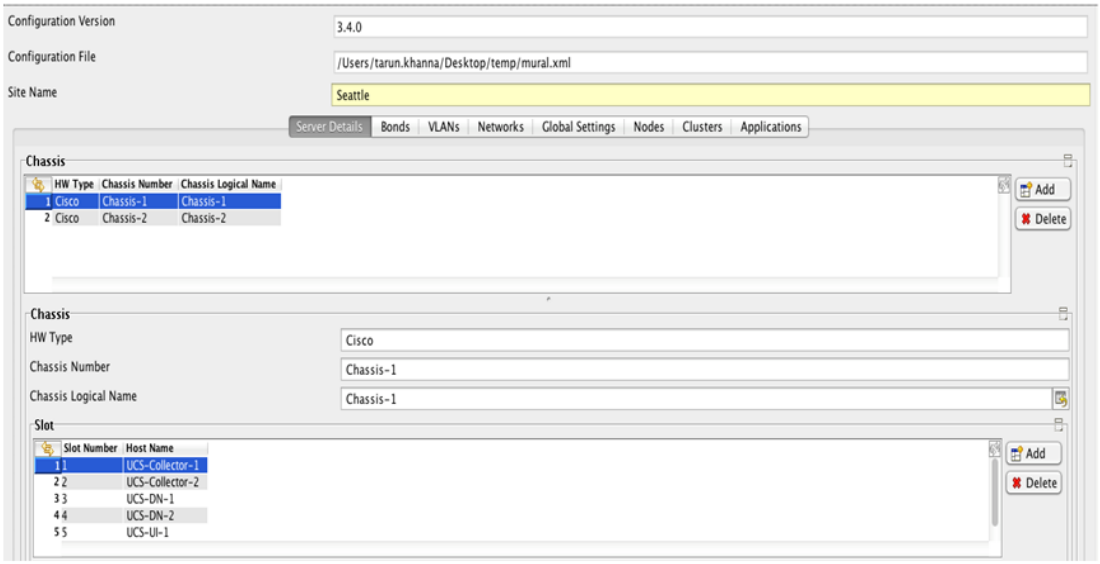
Verifying MAC Address and LUN WWIDs

The images used in this section show an example configuration. Your local configuration may vary. It is important to verify in the GMS interface that the configuration for all nodes in your deployments have been configured.

To verify the configurations for MAC addresses of interface bindings and storage partitions for each node:

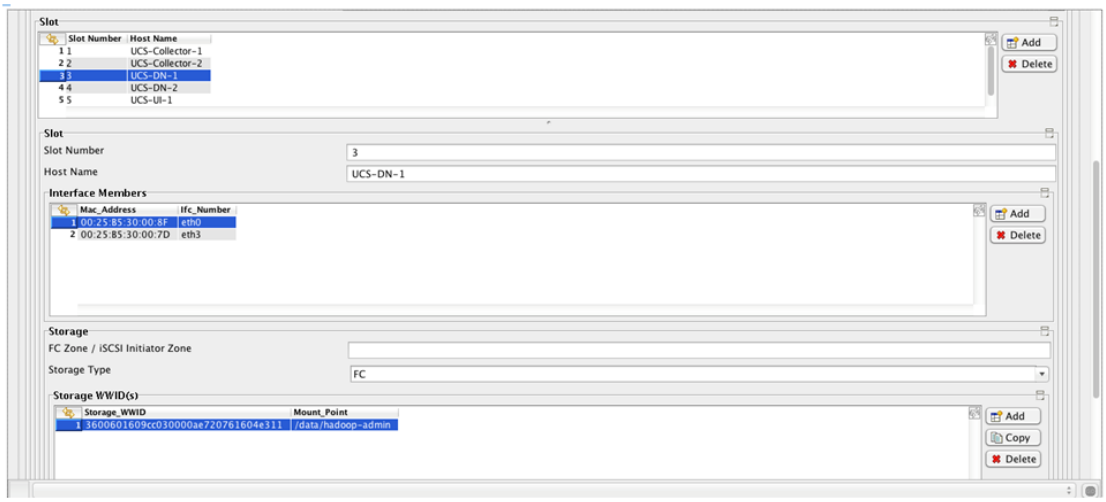
1. Select the chassis. The hardware type, chassis number and chassis logical name are displayed, as well as the list of nodes (blades) in each slot:

 Copyright © 2014, Cisco Systems, Inc.



2. Select the node.

3. Under **Interface Members**, verify the MAC address of the interface binding for the node against the information recorded for your system in your CIQ:



4. Under **Storage**, verify the Storage WWIDs (see previous figure). On the initial configuration of the deployment topology, EMC is not yet configured and you must enter dummy values. After EMC is configured, you will return to this procedure and enter the actual values for the WWIDs of the LUNs.

Tip: Even if entering dummy values for the WWIDs, you must make sure that they are unique among all the nodes (with the exception of the WWIDs that are shared between Insta1 and Insta2). If the WWIDs are not unique, the validation step at the end of the procedure fails with error message and you cannot save the configuration file until it is corrected.

5. If you need to change any settings, click **Add**, create a new entry, and then select the old entry and click **Delete**.
6. Repeat steps 1 through 3 for each of the following nodes in Chassis 1:
   - First Collector node
   - Second Collector node
   - First Compute node
   - Second Compute node
   - First UI/Caching (Rubix) node
   - Second UI/Caching (Rubix) node
   - First Insta node
   - Second Insta node

     Tip: Ensure that the same WWID is assigned to **dbroot1** for both Insta 1 and Insta 2 in the GMS configuration, and that the WWID for the **pgsql** storage partition is unique for each Insta node.
7. Repeat steps 1 through 4 for each of the nodes for Chassis 2.

Verifying Global Settings

Verify the global settings for the system (DNS, default gateway, and NTP), and IP addresses configured for each node cluster and for the standalone GMS node.

The Collector nodes have three IP addresses and all other nodes have two IP addresses:

- **Internal** or Control Network—The network that the nodes will use to communicate among themselves. For the **eth0** interface, we recommend you

 Copyright © 2014, Cisco Systems, Inc.

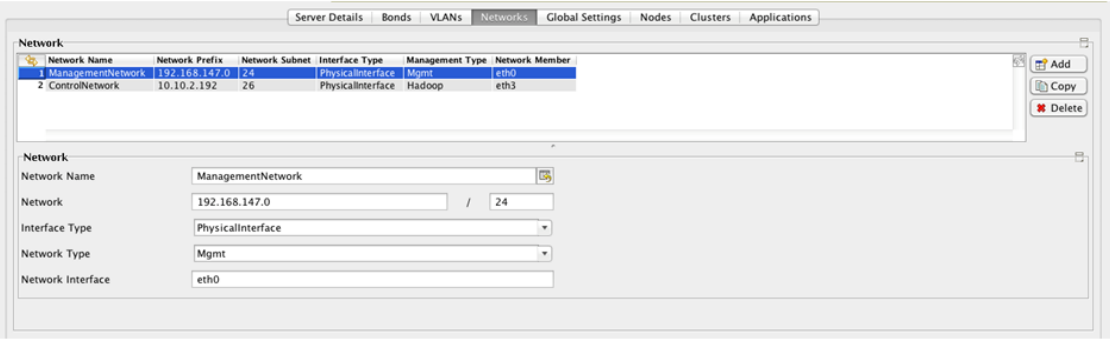
use a control IP address of the format 192.*x.x.x*.

- **External** or Management Network—The management network that will be used to SSH to the blades. For the **eth1** interface, we recommend you use a physical or management IP address of the format 10.*x.x.x*. In the example below, **eth3** is used for this interface. The management IP address can be used for the **eth0** interface when an internal network is not available.

To verify the network IP address, global settings, and IP addresses of the nodes:

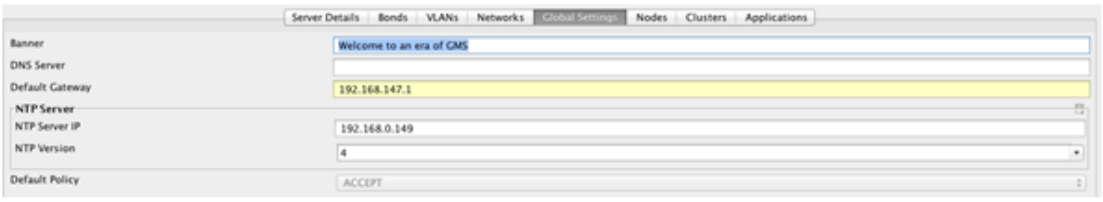
1. Under the **Networks** tab, verify the following fields for the internal (Management Network) and external (Control Network) networks:

    - Network name
    - Network prefix
    - Network subnet
    - Interface type
    - Management type
    - Network member

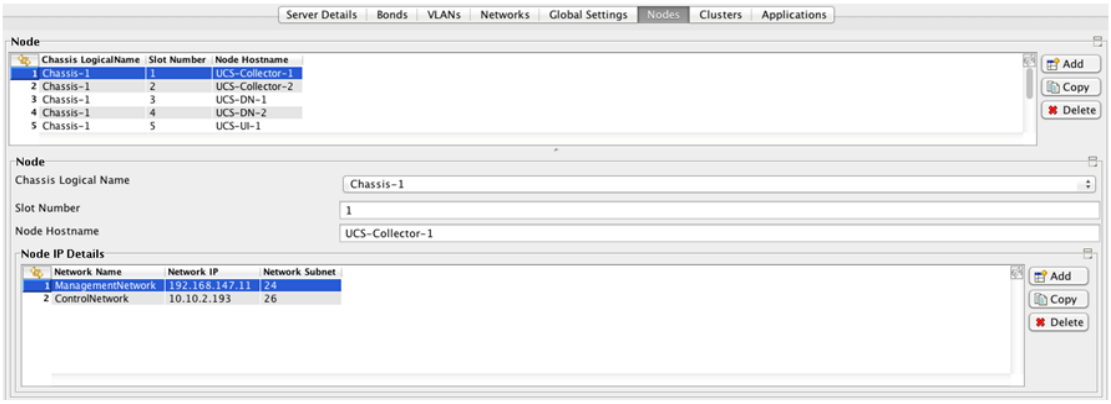


2. Under the **Global Settings** tab, enter the following information for the MURAL system:

    - DNS server
    - Default gateway
    - NTP server IP address
    - NTP version



3. Under the **Nodes** tab, define the internal and external IP addresses and subnet masks of the each node:



Configuring the Node Clusters

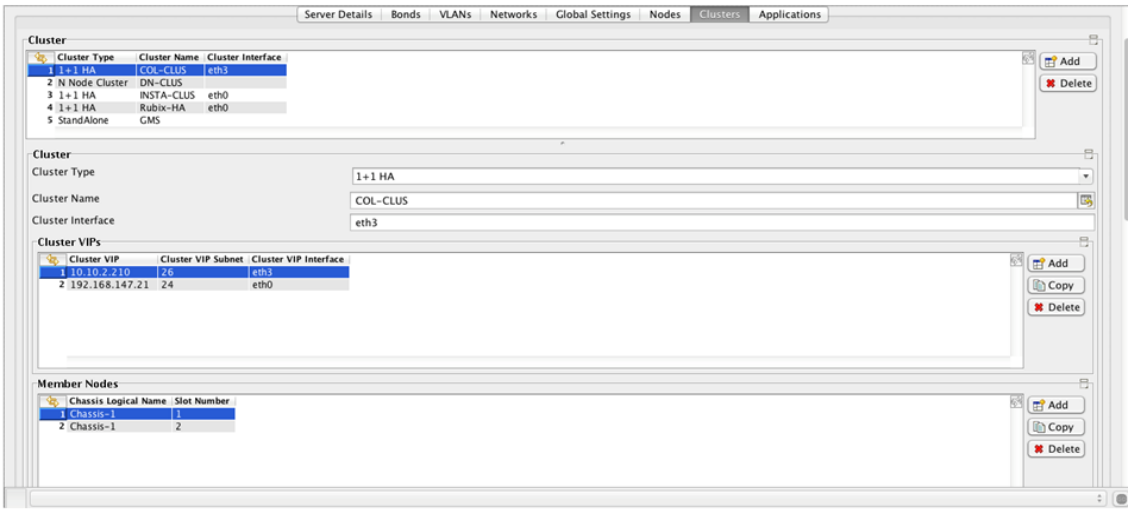
The **mural.xml** file included with your installation package includes configurations of the node clusters. The figures below show an example—the actual configurations for your setup might be different depending on how much traffic is being sent to the MURAL system.

During initial configuration, you must create clusters for each type of node and add the member nodes to each cluster.

To configure a cluster:

1. In the GMS interface, select the **Clusters** tab.
2. From the Cluster Type list, select the cluster you want to configure or verify.
3. Select a **Cluster Type** from the drop-down list, enter the **Cluster Name**, and **Cluster Interface**, and click **Add**. For example, the following figure shows a Collector cluster called COL-CLUS:

Copyright © 2014, Cisco Systems, Inc.



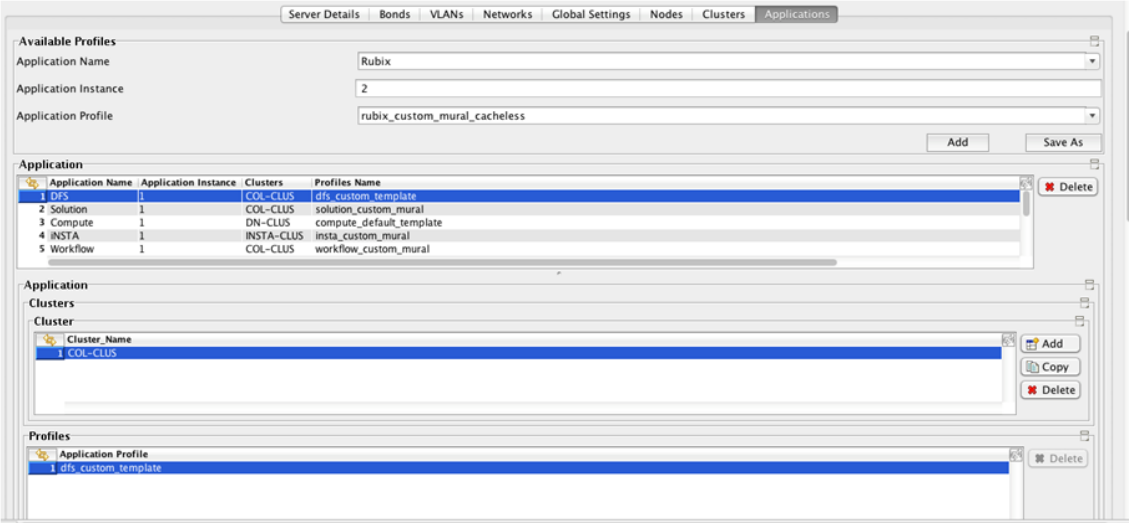
4. Select the newly created cluster, and under **Cluster VIPs**, click **Add** and enter the virtual IP address of the new cluster.
5. Repeat steps 2 through 4 for the following nodes:
    - Collector
    - Compute
    - Insta
    - Rubix
6. Define the GMS cluster related information. The GMS cluster type is **Stand-alone**.

Specifying Component Profiles and Application-Specific Settings

This section describes how to load profiles for each component. These profiles set the configuratons for the components. You can then change individual settings to adjust the configuration for your environment.

To load the profiles that you want to associate with each component:

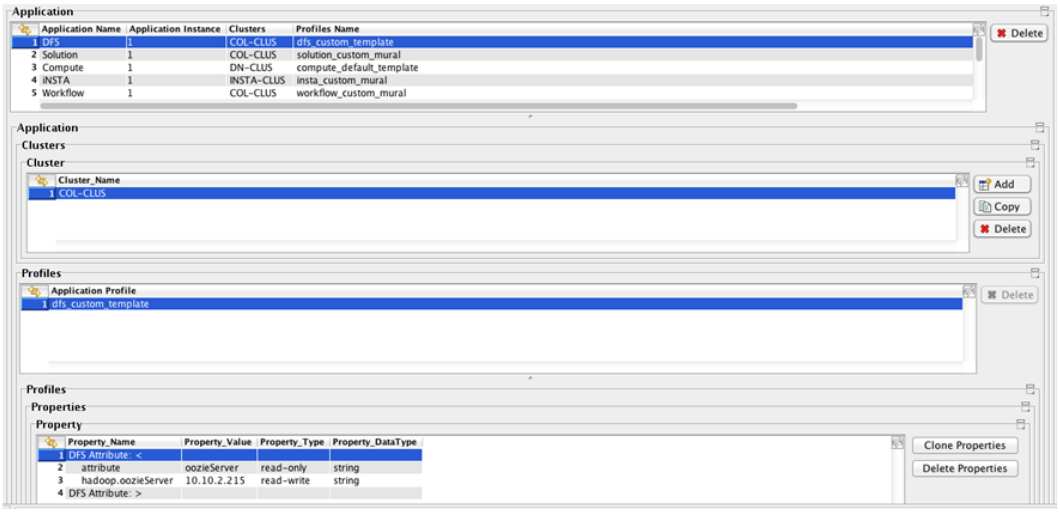
1. Select the **Applications** tab:



2. Under **Application Name**, select the name of the system component from the drop-down menu.
3. Under **Application**, the GMS interface displays the list of profiles. Select the profile that you want to attach to the system component.
4. Attach the DFS template to the Collector cluster. This template configures HDFS on the cluster.
    a. Under **Available Profiles** > **Application Name**, select **Collector**.
    b. Under **Application**, select **DFS**.
    c. Under **Application** > **Clusters**, select the Collector cluster and verify that the name is correct for your setup (in this example, **COL-CLUS**).
    d. In the DFS template, update Hadoop Oozie server to be control VIP of

Copyright © 2014, Cisco Systems, Inc.

the Collector blade, as shown below:

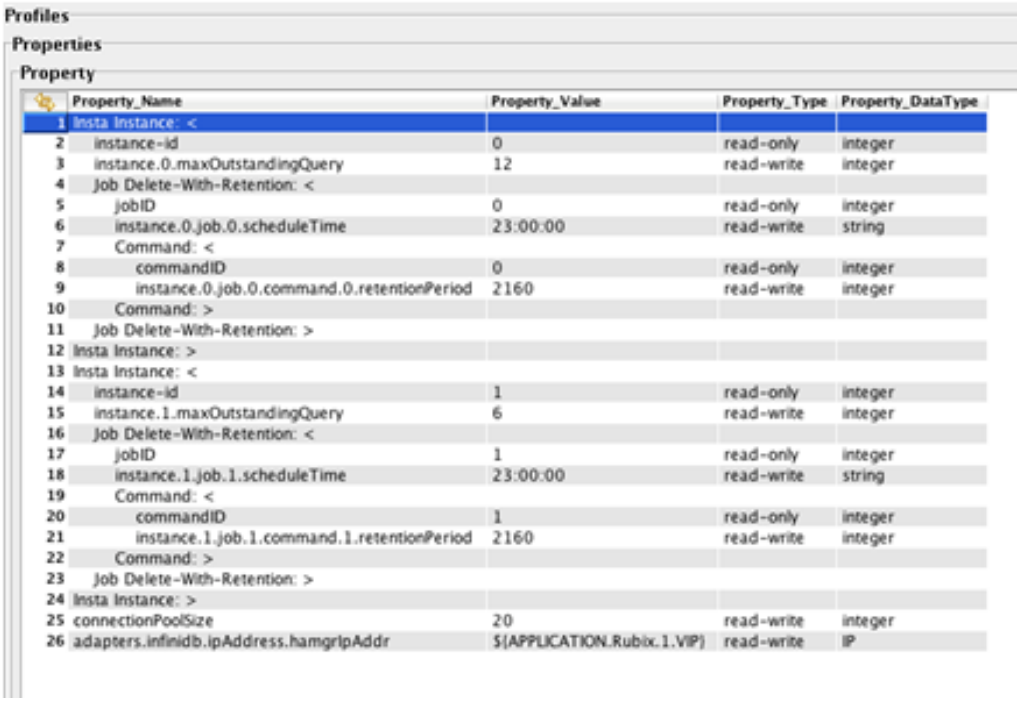


5. Attach the solution-based profile to the Collector cluster. Verify that the cluster name and the expected start date (**DataStartTime**) are correct for your setup.

   Note: It is safe to set the time here as GMS does not run all of the jobs. Later in this installation procedure you will change the Oozie start time and start all of the jobs (see "Set the Data Start Time" on page 123).

6. Attach the compute-based profile to the Compute cluster and verify that the name is correct for your setup.

7. Attach the Insta-related profile to the Insta cluster, and verify that the name of the cluster is correct for your setup. The following screen shows an example Insta profile.



Profiles

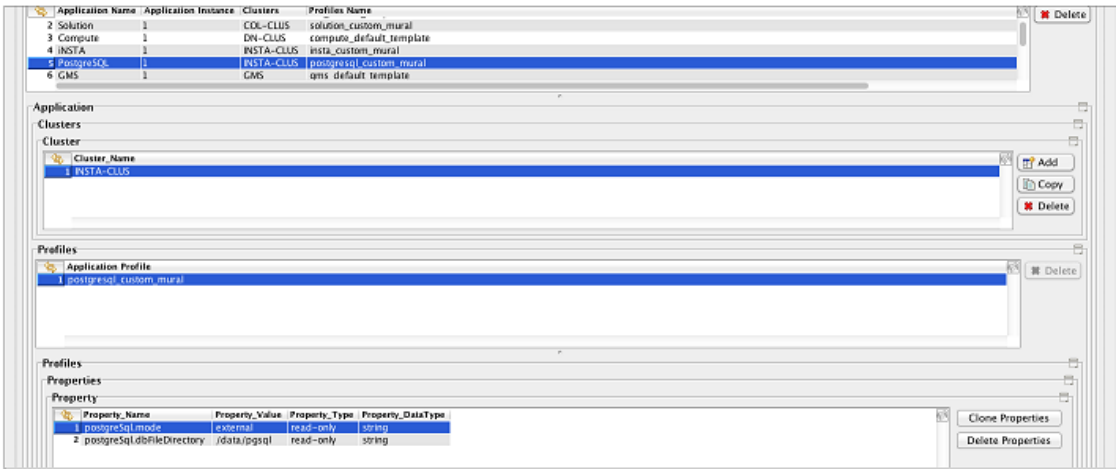
Properties

Property

	Property_Name	Property_Value	Property_Type	Property_DataType
1	Insta Instance: <			
2	instance-id	0	read-only	integer
3	instance.0.maxOutstandingQuery	12	read-write	integer
4	Job Delete-With-Retention: <			
5	jobID	0	read-only	integer
6	instance.0.job.0.scheduleTime	23:00:00	read-write	string
7	Command: <			
8	commandID	0	read-only	integer
9	instance.0.job.0.command.0.retentionPeriod	2160	read-write	integer
10	Command: >			
11	Job Delete-With-Retention: >			
12	Insta Instance: >			
13	Insta Instance: <			
14	instance-id	1	read-only	integer
15	instance.1.maxOutstandingQuery	6	read-write	integer
16	Job Delete-With-Retention: <			
17	jobID	1	read-only	integer
18	instance.1.job.1.scheduleTime	23:00:00	read-write	string
19	Command: <			
20	commandID	1	read-only	integer
21	instance.1.job.1.command.1.retentionPeriod	2160	read-write	integer
22	Command: >			
23	Job Delete-With-Retention: >			
24	Insta Instance: >			
25	connectionPoolSize	20	read-write	integer
26	adapters.infinidb.ipAddress.hamgrIpAddr	${APPLICATION.Rubix.1.VIP}	read-write	IP

8. Attach the workflow-related profile to the Collector cluster. The workflow profile contains configuration settings for the MapReduce jobs.

9. Attach the Postgres-related profile for the Postgres application and assign it to the Insta cluster, as shown in the following screen:



	Application Name	Application Instance	Clusters	Profiles Name
2	Solution	1	COL-CLUS	solution_custom_mural
3	Compute	1	DN-CLUS	compute_default_template
4	iNSTA	1	INSTA-CLUS	insta_custom_mural
5	PostgreSQL	1	INSTA-CLUS	postgresql_custom_mural
6	GMS	1	GMS	gms_default_template

Delete

Application

Clusters

Cluster

	Cluster_Name
1	INSTA-CLUS

Add

Copy

Delete

Profiles

	Application Profile
1	postgresql_custom_mural

Delete

Profiles

Properties

Property

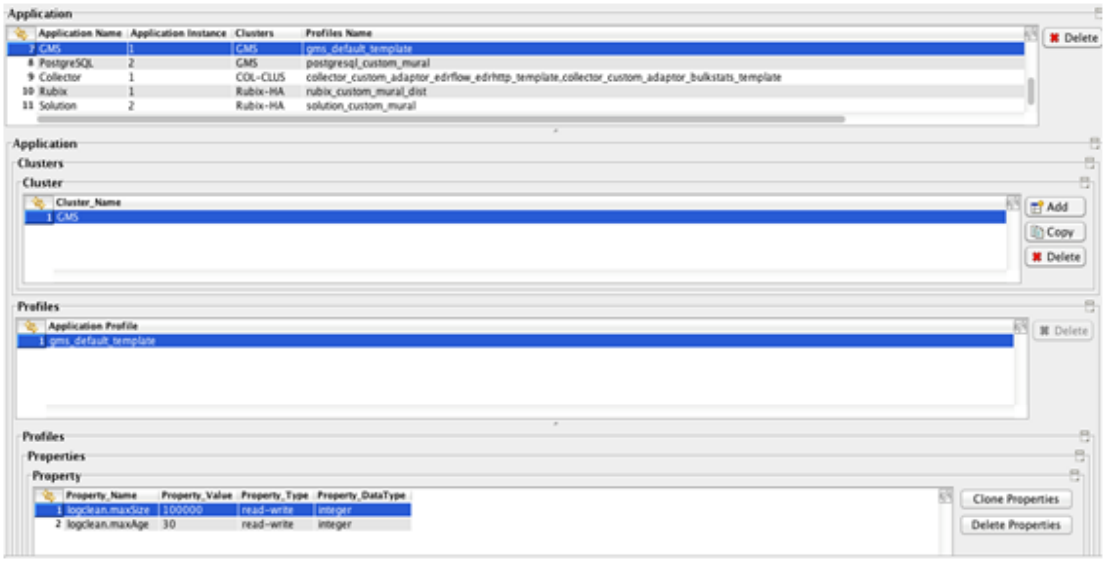
	Property_Name	Property_Value	Property_Type	Property_DataType
1	postgreSql.mode	external	read-only	string
2	postgreSql.dbFileDirectory	/data/pgsql	read-only	string

Clone Properties

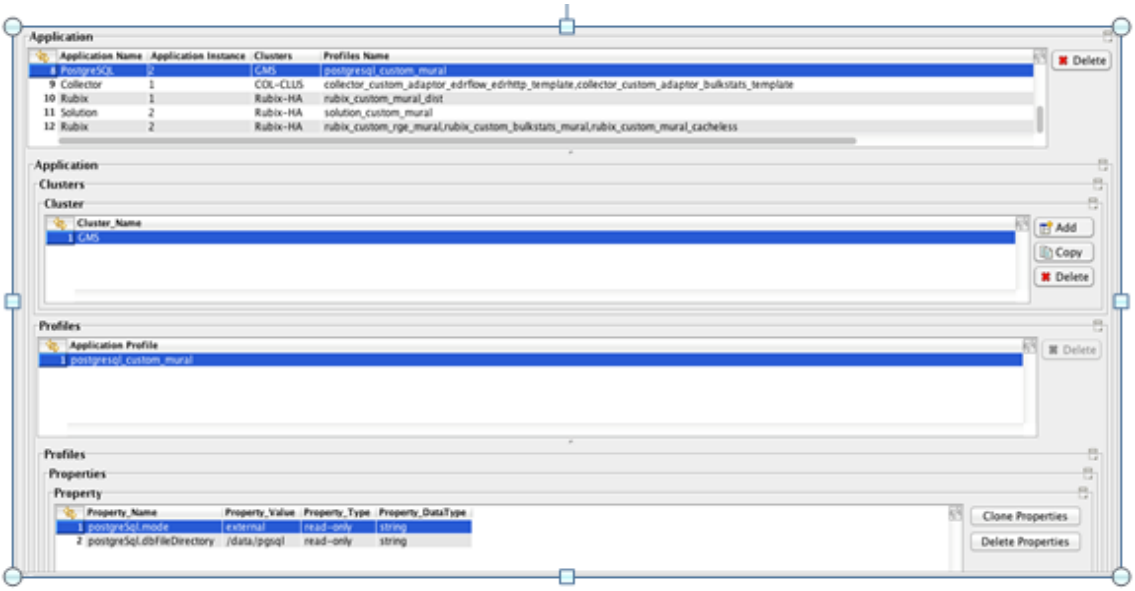
Delete Properties

10. Attach the GMS-related profile for the GMS application and assign it to the GMS cluster:

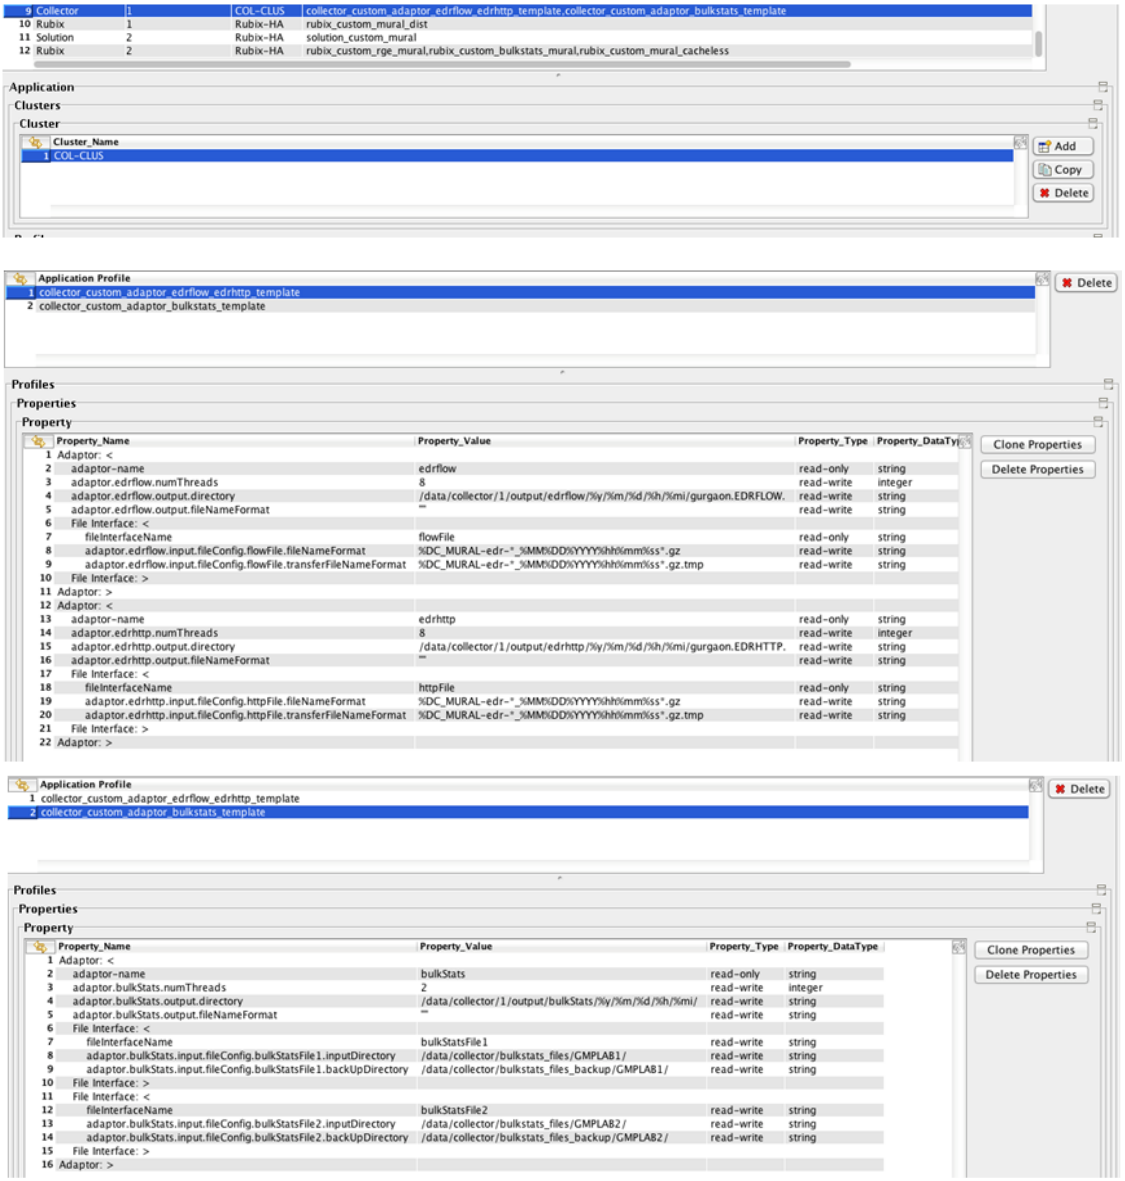
Copyright © 2014, Cisco Systems, Inc.



11. Attach the Postgres-related profile for the Postgres application and assign it to GMS cluster:



12. Attach the Collector-based templates (**collector_custom_adaptor_edrhttp_edrflow_template** and **collector_custom_adaptor_bulkstats_template**) to the Collector cluster, specifying **Collector** as the Application Name:



Verify that the following parameters are correct for your setup:

- Cluster name, specified by the parameter `Cluster_Name`.
- BulkStats file input directory, specified by the parameter `adaptor.bulkStats.input.fileConfig.bulkStatsFile1.input-Directory`.

  The format of the value is `/data/collector/bulkstats_files/`*`gateway_name`*.

  An example value is `/data/collector/bulkstats_files/GMPLAB1/`,

Copyright © 2014, Cisco Systems, Inc.

where the ASR is located in the DC called GMPLAB1.

- BulkStats filebackup directory, specified by the parameter `adaptor.bulkStats.input.fileConfig.bulkStatsFile1.backupDirectory`.
  The format of the value is `/data/collector/bulkstats_files/`*`gateway-name`*.
- BulkStats file name format, specified by the parameter `adaptor.bulkStats.input.fileConfig.bulkStatsFile1.fileNameFormat`.
  The format of the value is `*_%YYYY%MM%DD%hh%mm%ss`.
- Output directory for EDR flow, specified by the parameter `adaptor.edrflow.output.directory`.
  The format of the value is `/data/collector/1/output/edrflow/`*`%y/%m/%d/%h/%mi/gateway-name`*.
- Output directory for EDR HTTP, specified by the parameter `adaptor.edrhttp.output.directory`.
  The format of the value is `/data/collector/1/output/edrhttp/`*`%y/%m/%d/%h/%mi/gateway-name`* (where `%y/%m/%d/%h/%mi/` indicates the format of the time value).

13. Attach the appropriate profile to the Rubix cluster and verify the values of the following parameters:
    - `application.atlas.sessionCookieDomain`
      For example, if the URL to be used is **https://today.mural.com** then the correct value is `mural.com`.
    - `application.atlas.timeZone`
      The time zone of location of the ASR, for example, `America/Argentina/Buenos_Aires`.

- `application.bulkstats.sessionCookieDomain`
- `application.reportAtlas.sessionCookieDomain`
- `application.reportAtlas.rgeUrl`

  URL to be used, for example, `https://today.mural.com:30443`.
  Note that the port number must be 30443.
- `application.reportAtlas.bulkStatsURL`
- `application.rge.mailNotificationHost`

  Customer mail host, for example, `mx1.cisco.com`.
- `application.rge.mailNotificationSender`

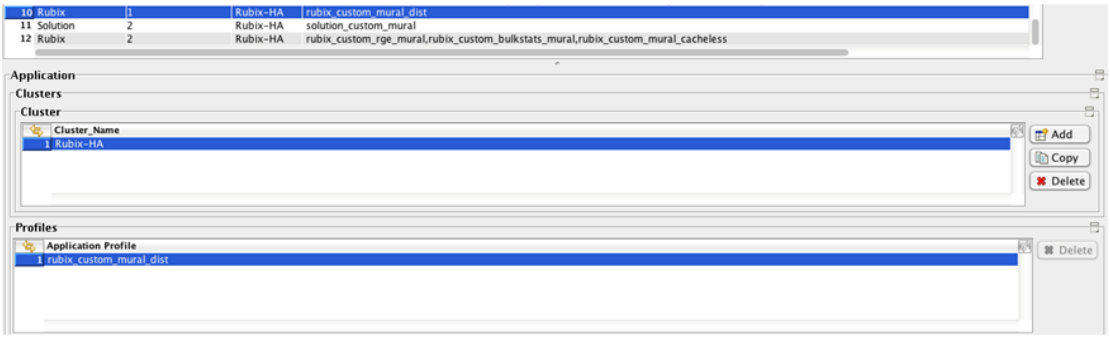
  Destination email address, for example, `support@cisco.com`.
- `application.rge.rgekeyStoreAlias`

  Caution: Do not change this value from `app2.guavus.com`.
- `application.bulkstats.timeZone`

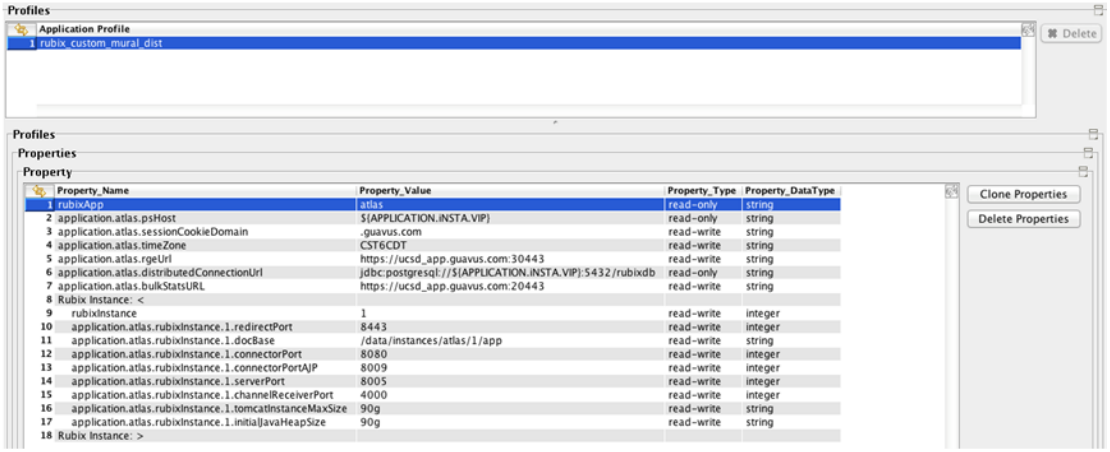
  Time zone of the location of the ASR.
- `application.reportAtlas.timeZone`

  Time zone of the location of the ASR.

The following screens show an example:



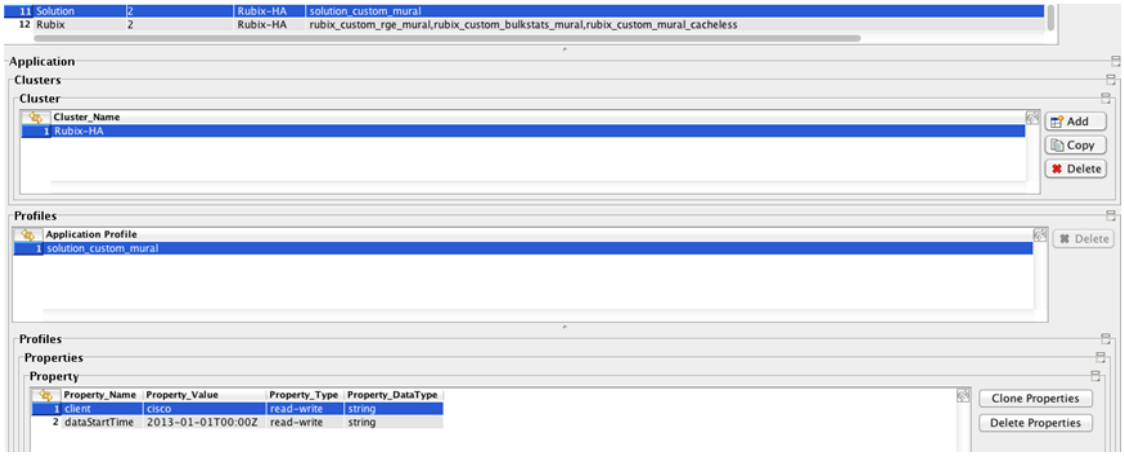
Copyright © 2014, Cisco Systems, Inc.



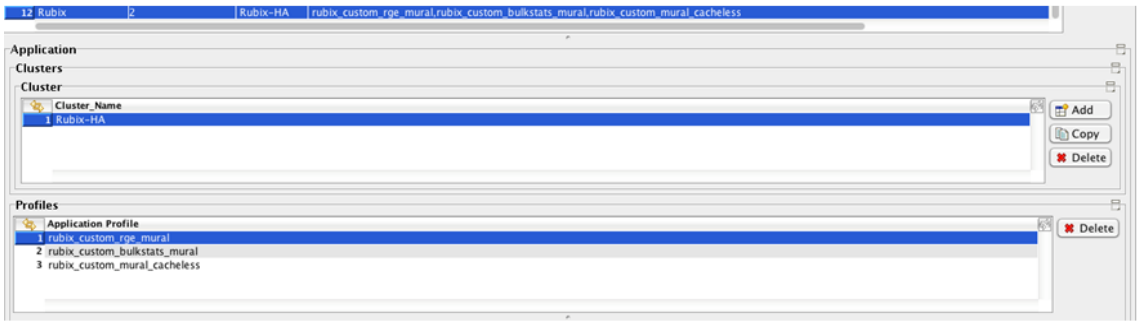
Note: Set the **application.timeZone** property to the same **TimeZone** value for all applications (such as Atlas, RGE, and Bulkstats).

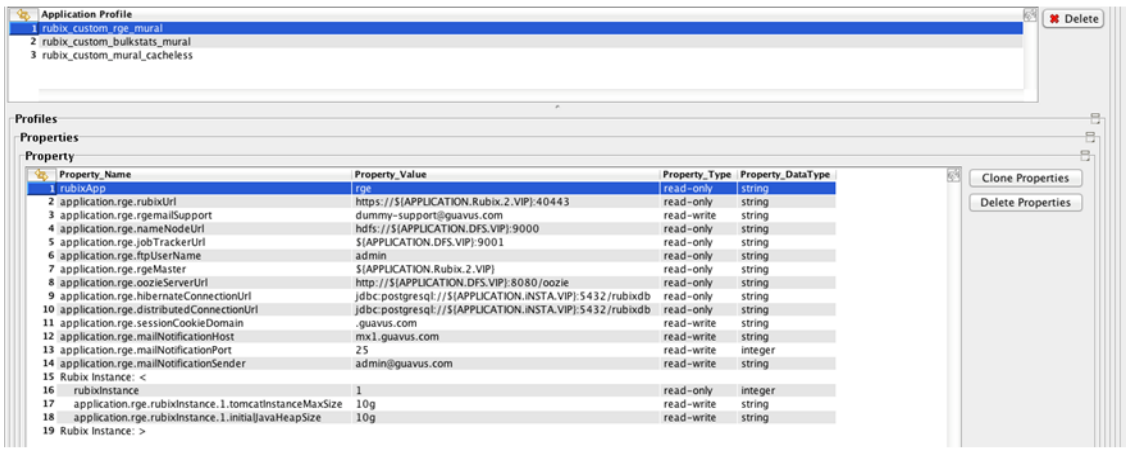
TimeZone String	Actual Name
CST	United States central standard time
EST	United States eastern standard time
Chile/Continental	Chile time
America/Argentina/Buenos_Aires	Argentina time
CET	Central European time
GMT	Greenwich mean time

14. Specify the solution-based template to the Solution application and attach it to the Rubix Cluster as Instance 2:

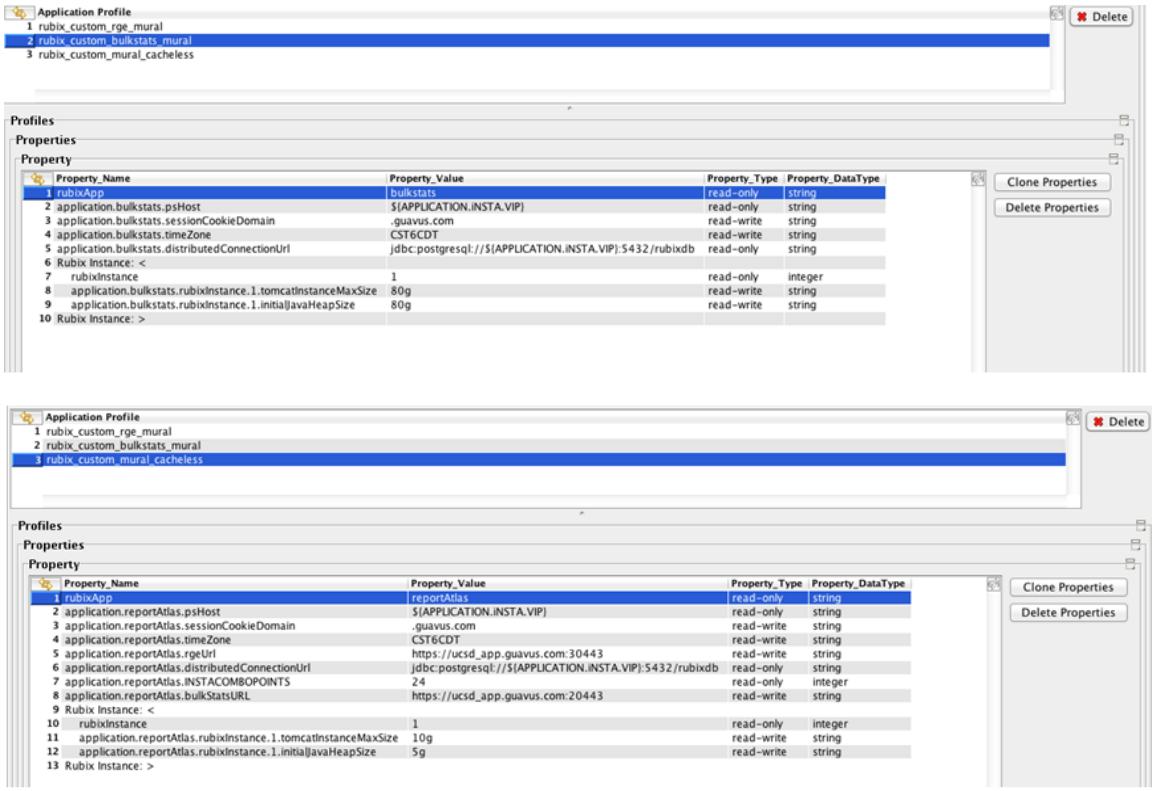


15. Ensure that Application Profiles are attached to the Rubix cluster as in the following sample screen shots:





Copyright © 2014, Cisco Systems, Inc.



Validating the Configuration

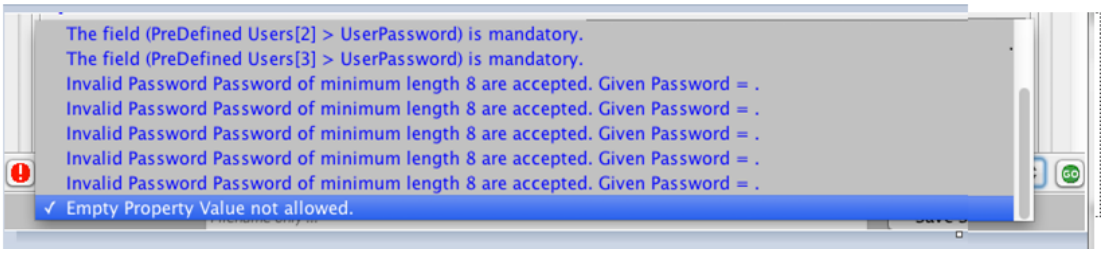
After completing all configurations on the **Server Details**, **Networks**, **Global Settings**, **Nodes**, **Clusters**, and **Applications** tabs as described in the previous sections, you must validate the configuration settings.

To validate the GMS configuration:

1. Click **Validate** on the bottom bar of the GMS interface.

   If any field or value entered is incorrect, a list of errors is displayed.

2. For each error in turn, click the error and then the **Go** button to display the entry in the GMS configuration that is causing the error.



3. Correct any errors.
4. Click **Validate** again. When all errors have been resolved, the message **Validation successful** displays on the screen.
5. Click **Save Server File** to save the configuration to the GMS server.

   Note: The **Save Server File** is not operational until all validation errors have been corrected.To save the file without correcting all errors and completing the validation, click **Save Local File**.

To activate the configuration after you have made changes, you must run the `gms config mural.xml activate` command. Until you activate the new script, the system runs the previously saved configuration.

Copyright © 2014, Cisco Systems, Inc.

Manufacturing All Other Blades

To manufacture the blades for all nodes other than the GMS node:

1. SSH to the GMS server using the management IP address and start the installation on all UCS nodes:

   ```
   > en
   # conf t
   (config) # gms config mural.xml activate
   ```

   Notes:
   The **mural.xml** file is the same file used during the GMS configuration.

   If you make any changes to the **mural.xml** file, for example during troubleshooting, you must run this command again to activate the changed file.

2. Manufacture all the blades using PXE boot via GMS:

   ```
   > en
   # conf t
   (config) # image fetch scp://admin@ IP-local-remote-machine/
   directory-path-on-localmachine/iso-image-name
   (config) # image mount iso-image-name
   ```

   For example:

   ```
    (config) # image fetch
    scp://admin@50.204.88.45/data/ftp/software/images/cisco_mural/
   mural3_2mr1/
    mfgcd-guavus-x86_64-20130925-174106.iso

    (config) # image mount mfgcd-guavus-x86_64-20130925-174106.iso
   ```

   Note: *IP-local-remote-machine* is the IP address of the machine where the ISO image for the build has been downloaded.

   The above command will show output similar to the following:

```
guavus-3000ae (config) # image mount
mfgcd-guavus-x86_64-20130827-085941.iso
 Copying linux...
 Copying rootflop.img...
 Copying image.img...
```

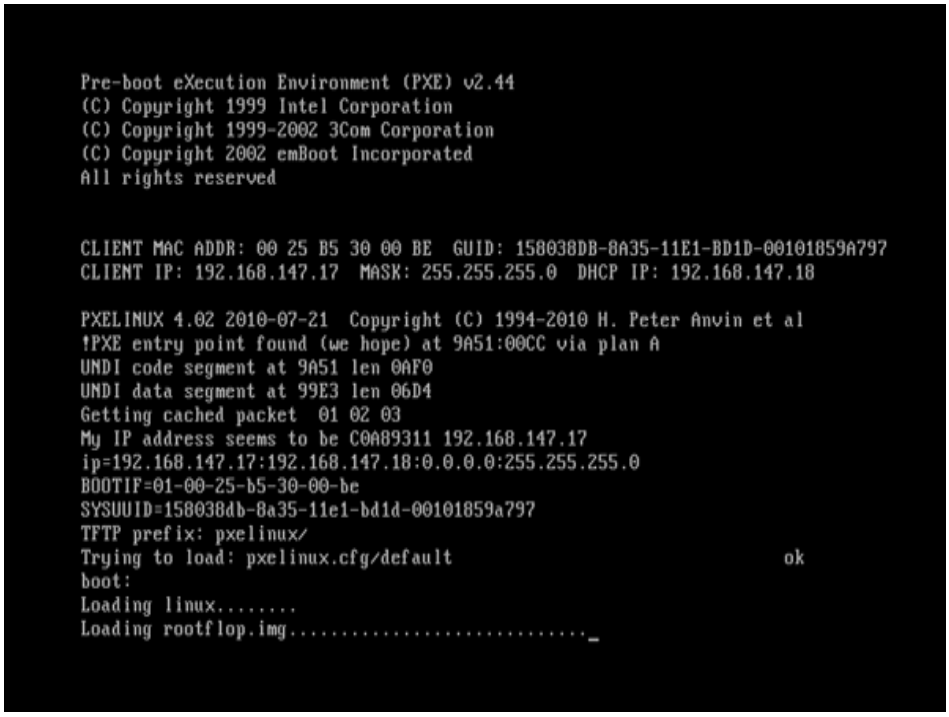
3. Verify the image on the GMS node:

```
# ls /var/opt/tms/images/
mfgcd-guavus-x86_64-20130925 174106.iso
```

4. Complete the following steps to reboot each blade from the KVM Console:
    a. Open the KVM Console of the blade.
    b. Press **CTRL-ALT-DEL**, as explained in "Manufacturing the GMS Blade" to reboot the blade. You can also click **Reset** from the top of the KVM console.
    c. After the prompt, press **F12** as soon as possible to boot from the network.
    d. Repeat steps a through c for all blades, except those designated for GMS.

 Copyright © 2014, Cisco Systems, Inc.

The following figure shows the results of clicking **F12**.



Once the blades start booting from the network, GMS pushes the image on all the blades using PXE boot for the manufacture process to start on each blade in parallel.

A blade takes approximately 20 minutes to manufacture with the new image. Wait until the last blade for which PXE boot was issued has been manufactured. A login prompt is displayed once the image has been manufactured on a blade.

Verify Zoning/Flogi on the Fabric Interconnect

1. Log on to the Fabric Interconnect (FI) using SSH.
2. Run the `connect nxos A` command to connect to the Fabric Interconnect.

   Note: Ensure the FCID for all of the blades is shown. The command can be changed to print the active zoneset (for example, to print for **FI A**).
3. Run the **show zoneset active** command to print, as shown for FI A and FI B in this example:

Zoneset Active for FI A

```
mural-norse-lab-A(nxos)# show zoneset active

zoneset name ucs-mural-norse-lab-vsan-3010-zoneset vsan 3010
  zone name ucs_mural-norse-lab_A_12_UI1_vHBA-A vsan 3010
  * fcid 0x100003 [pwwn 20:00:00:05:ad:1e:11:df]
  * fcid 0x1000ef [pwwn 50:06:01:60:3e:a0:28:d2]
  * fcid 0x1001ef [pwwn 50:06:01:69:3e:a0:28:d2]

  zone name ucs_mural-norse-lab_A_11_UI2_vHBA-A vsan 3010
  * fcid 0x100006 [pwwn 20:00:00:05:ad:1e:11:ff]
  * fcid 0x1000ef [pwwn 50:06:01:60:3e:a0:28:d2]
  * fcid 0x1001ef [pwwn 50:06:01:69:3e:a0:28:d2]
```

Zonset Active for FI B

```
mural-norse-lab-B(nxos)# show zoneset active

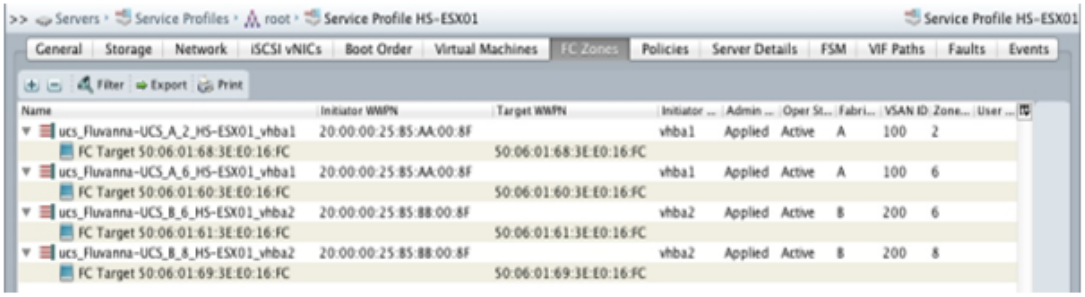
zoneset name ucs-mural-norse-lab-vsan-3020-zoneset vsan 3020
  zone name ucs_mural-norse-lab_B_24_UI1_vHBA-B vsan 3020
  * fcid 0x420007 [pwwn 20:00:00:05:ad:1e:11:2e]
  * fcid 0x4200ef [pwwn 50:06:01:61:3e:a0:28:d2]
  * fcid 0x4201ef [pwwn 50:06:01:68:3e:a0:28:d2]

  zone name ucs_mural-norse-lab_B_23_UI2_vHBA-B vsan 3020
  * fcid 0x420009 [pwwn 20:00:00:05:ad:1e:11:5e]
  * fcid 0x4200ef [pwwn 50:06:01:61:3e:a0:28:d2]
  * fcid 0x4201ef [pwwn 50:06:01:68:3e:a0:28:d2]
```

NOTE: Verify that the output reports the FCID for all WWPNs and hosts.

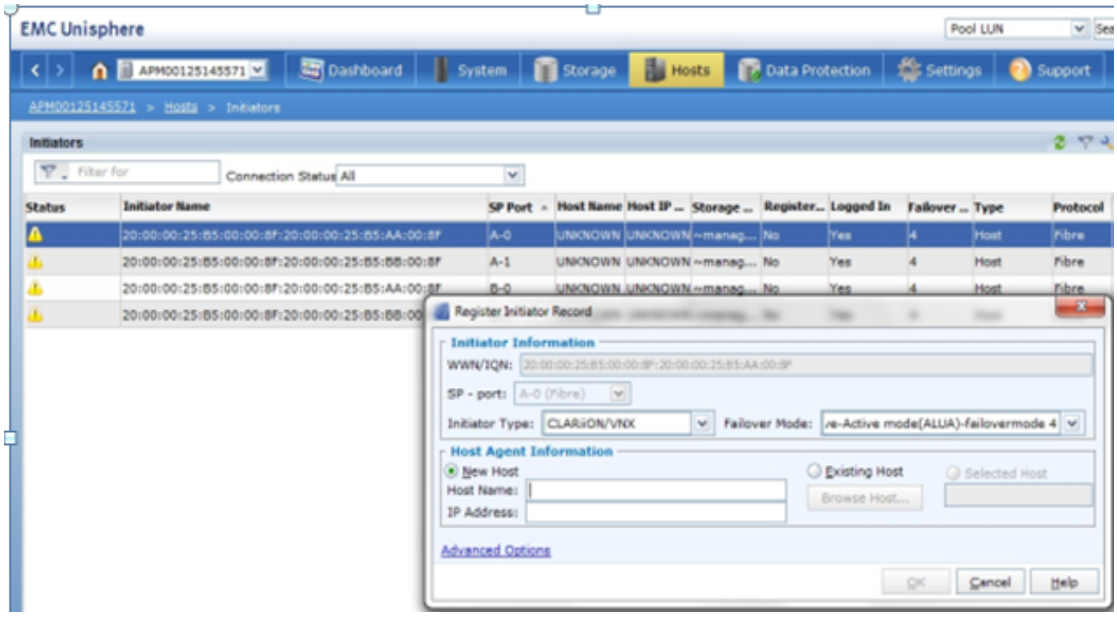
Use UCS Manager to verify the zoning

Ensure that the initiator and target WWPNs are in the same zone, as shown in the following figure:

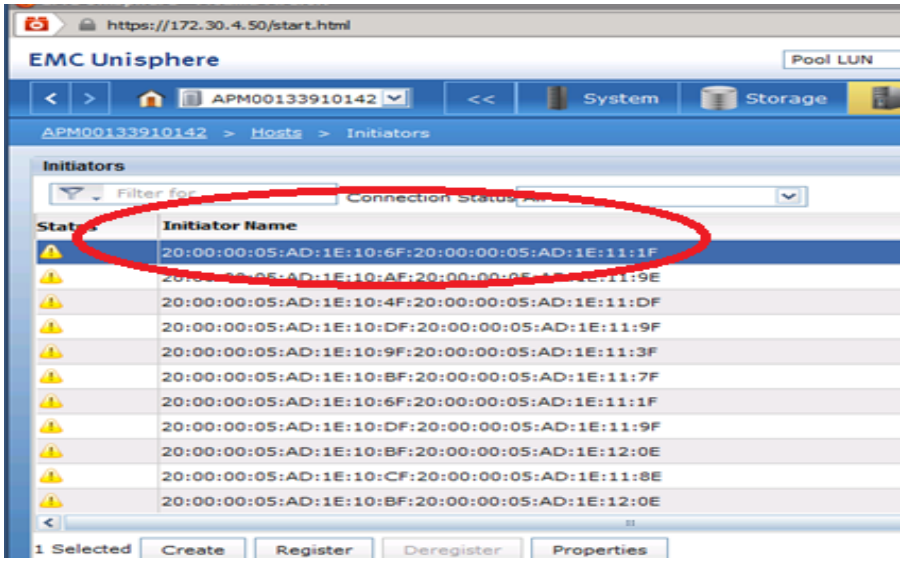


 Copyright © 2014, Cisco Systems, Inc.

Hosts are ready to be registered, as in the following example, which shows hosts on EMC:



NOTE: WWPNs on the UCS Manager are on **SAN > Pools > WWPN Pools > WWPN Pool** (the ID used in this installation is **wwpn-ppol1/Initiator**). On the EMC in the following figure, the Initiator Name format is WWNNWWPN (where WWNN is first, then WWPN):



Setting Up Data Storage On the EMC

This section describes how to allocate data storage to each node in the MURAL system. EMC is used to manage the storage environment for MURAL. Storage includes the HDFS connected to the Collector and Compute nodes and the columnar database used by the Insta nodes.

EMC Hardware Installation Prerequisites

Before beginning, verify that you have completed the following EMC hardware installation tasks:

- Installed the VNX storage chassis and standby power supply (SPS) chassis in the rack following the instructions provided in the EMC Unisphere installation guide (EMC P/N 300-012-924) included with the hardware.
- Connected the SPS to the storage processor (SP) management ports using the cables provided with the product and following the instructions provided in the EMC Unisphere installation guide.
- Connected the power cords from SPS A / SPS B to SP A and SP B, and from SPS A and SPS B to PDUs following the instructions provided in the EMC Unisphere installation guide.
- Installed the Fibre Channel SFP+ transceiver, included with the hardware, in Ports 4 and 5 of both SP A and SP B.

Note: Leave the storage system uncabled to the server array until after initialization is complete.

Also obtain the IP addresses for the two SP management ports and other items listed in the following table. Refer to your CIQ to obtain the values.

Item	Value
SP A management port IP	
SP B management port IP	
Subnet mask and gateway for above	
Admin name/password	

Copyright © 2014, Cisco Systems, Inc.

Item	Value
Storage system serial number	
Scope	
DNS server address (optional)	
Time server address	
In-bound email address	

Note: The following IP addresses cannot be used128.121.1.248 through 128.121.1.56, 192.168.1.1, and 192.168.1.2.

Configuring Your Base IP to Access the EMC System

The default IP adresses for the EMC system are **1.1.1.1** and **1.1.1.2**. By configuring your laptop to a similar range, you can connect to **1.1.1.1** using a web browser and reconfigure the IP address information:

1. Configure your laptop's IP adress to **1.1.1.4/24**,
2. Connect a cable to Service Processor A.
3. Open a web browser to **http://1.1.1.1/setup**.
4. Reconfigure the IP adresses to the desired range.

Note: If you need to manually restart EMC during the set-up procedures, open a web browser to **http://1.1.1.1/setup**, log in as admin, and select the **restart** option.

Registering All MURAL Nodes on the EMC

From the EMC, you must associate each MURAL node with the WWPN of the correct fiber channel port on the storage arrays. The WWNN (world wide name node) is a unique identifier assigned to a port in a Fiber Channel fabric, in this case, each MURAL node. The world wide port name (WWPN) is the unique identifier for a fibre channel port, in this case on the storage arrays.

You must allocate memory to each node in the MURAL system. To do this, from the EMC Unisphere interface, define each WWPN node exactly as it is defined in

the UCS. If your EMC has been previously configured with MURAL, you still need to verify that the existing definition in EMC exactly matches the definition in the UCS.

Note: Alerts might be generated during the setup process indicating that nodes are not registered (Alert 0x721c). These are normal until provisioning is complete, and you only need to investigate those that occur after provisioning is complete.

Before You Begin

Before you begin to set up the EMC, gather the information about the nodes on the UCS from the UCS Manager and fill in the table below. Typically the last two digits are enough to identify the node on the EMC interface.

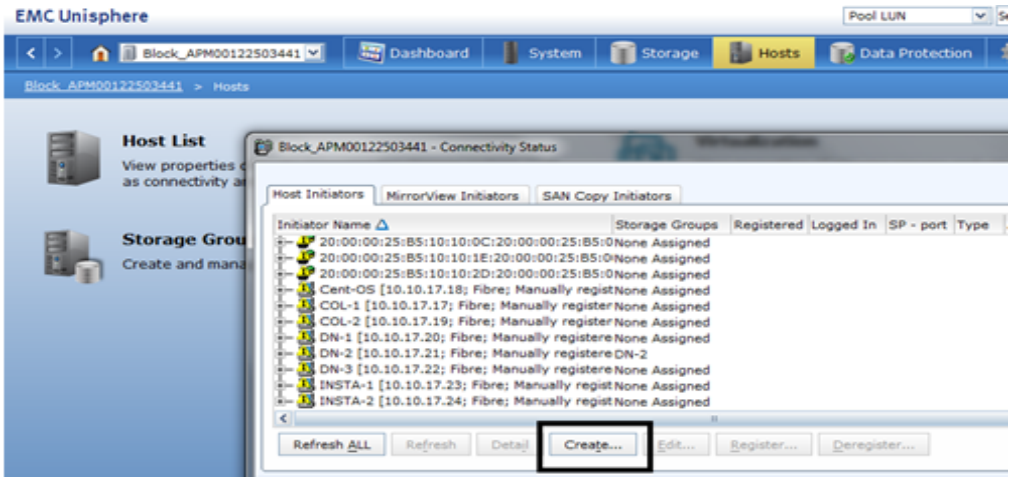
Node - WWNN	WWPN - 1	WWPN - 2
Collector1		
Collector2		
Compute1		
Compute2		
Compute3		
Compute4		
Insta1		
Insta2		
UI1		
UI2		

To register the MURAL nodes on the EMC:

1. On the EMC, mouse over (not click) **Hosts**, and select **Connectivity Status**.

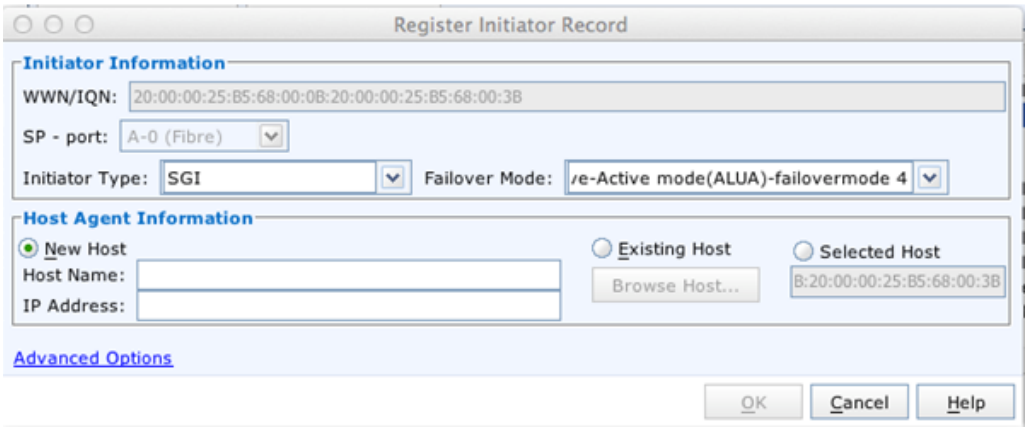
   The following screen displays:

Copyright © 2014, Cisco Systems, Inc.



2. If the correct WWPN is shown for a node, then the storage unit is already configured, and you just need to verify.

   If the node is not yet configured, click **Create**. The following screen displays:



3. The WWN for the FC port of the storage partition is shown. The WWN is a unique 16-digit hexadecimal number, such as 20:00:00:25:B5:68:00:08. This identifier is hard-coded into every FC host bus adapter (HBA). WWNs are important when setting up a storage area network (SAN)—each device has to be registered with the SAN by its WWN before the SAN will recognize it.

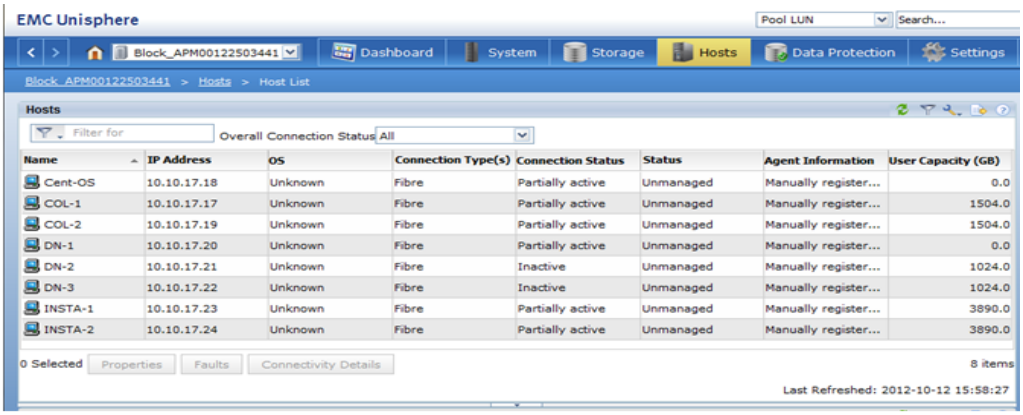
   Note: The WWN/IQN shown above is the same Initiator Name that you would find from the UCS Manager, as follows:

- From the UCS Manager, go to **SAN** > **Pools** > **wwnn-pool-wwnn-pool1**. The **Initiator** tab shows the WWNN as 20:00:00:05:AD:1E:10:6F. Note the related node from the **Assigned to** column. **Wwnn-pool1** is the name used during initial configuration using the UCS configuration script.
- Go to **SAN** > **Pools** > **wwpn-pool-wwpn-pool1**. The **Initiator** tab shows the WWPN for the same node (that is, GMS1) as 20:00:00:05:AD:1E:11:1F. Again **wwpn-pool1** is the name used during initial configuration using the UCS configuration script.

  In other words, on the EMC WWNN and WWPN are shown as WWNNWWPN for the node.

4. On the Register Initiator Record screen (shown above), do the following:
   - Enter the hostname and the IP address for the MURAL node.
   - Verify that the Initiator type shown in the pull-down list is **SGI**.
   - Verify that failover is set to mode 4 ALUA (the default setting).
5. Repeat Steps 1 through 5 for each of the nodes.
6. When you have registered all of the nodes, verify they are correctly configured—from the UniSphere management agent, go to **Dashboard** > **Hosts** > **Host List**.

   A list of hosts similar to the following is displayed:



Copyright © 2014, Cisco Systems, Inc.

Creating and Assigning RAID Groups and LUNs

First create RAID groups, then create LUNs and assign them to the appropriate RAID group.

Create RAID groups 5, 10, and 100, with the parameters specified in the following table:

RAID Group	RAID Type	Storage Pool ID	Disks
5	5	5	4,5,6,7,8
10	10	10	10,11,12,13
100 (not used)	unbound	100	0,1,2,3

Add any spare disks to RAID Group 5 or 100 depending on the hardware.

Tip: Prior to setting up the RAID groups, it might be useful to prepare a table similar to the example in "Setting Up Data Storage On the EMC" in order to ensure all disks required have been allocated.

After creating the RAID groups, create the LUNs. The following table shows example disk sizes. These sizes are examples only, and you should confirm the sizing for your specific environment before you begin.

RAID	RAID Group Name	LUN Name	LUN ID	Disk Size (GB)	Controller	Storage Pool	Host - MAP
10	RAID Group 10	INSTA-1	0	1945	FAB-A	INSTA-STR-1	INSTA NODE-1
10	RAID Group 10	INSTA-1-PGSQL	1	50	FAB-A	INSTA-STR-2	INSTA NODE-1
10	RAID Group 10	INSTA-2	2	1945	FAB-B	INSTA-STR-1	INSTA NODE-2
10	RAID Group 10	INSTA-2-PGSQL	3	50	FAB-B	INSTA-STR-3	INSTA NODE-2
5	RAID Group 5	COL-1	4	1024	FAB-A	COL-STR-1	COL NODE-1

RAID	RAID Group Name	LUN Name	LUN ID	Disk Size (GB)	Controller	Storage Pool	Host - MAP
5	RAID Group 5	COL-2	5	1024	FAB-B	COL-STR-1	COL NODE-2
5	RAID Group 5	DN-1	6	1024	FAB-A	DN-STR-1	DN-1
5	RAID Group 5	DN-2	7	1024	FAB-B	DN-STR-2	DN-2
5	RAID Group 5	DN-3	8	1024	FAB-A	DN-STR-3	DN-3
5	RAID Group 5	DN-4	9	1024	FAB-B	DN-STR-4	DN-4
5	RAID Group 5	UI-1	10	1024	FAB-A	UI-STR-1	UI-1
5	RAID Group 5	UI-2	11	1024	FAB-B	UI-STR-2	UI-2
5	RAID Group 5	GMS-1	12	TBD*	FAB-A	GMS-STR-1	GMS-1
5	RAID Group 5	GMS-2	13	TBD*	FAB-B	GMS-STR-2	GMS-2

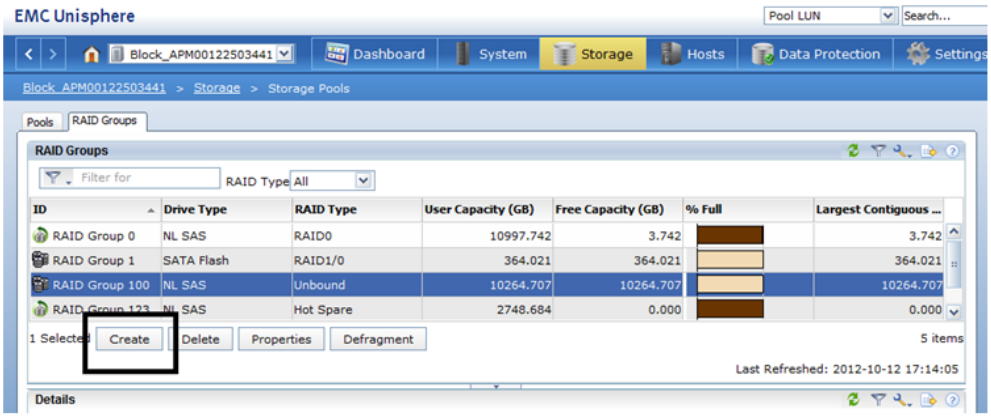
Note:

- Actual disk sizes will vary depending on the data throughput supported on the system.
- For the Insta-1-PGSQL and Insta-2-PGSQL disks, 50 GB is the minimum for a production setup. For any variations in size, such as for a lab environment, contact Technical Support.
- *Be sure to create the LUNs for the GMS nodes. Contact Technical Support for the correct size of the GMS LUNs for your environment.

To create and assign RAID groups and LUNs:

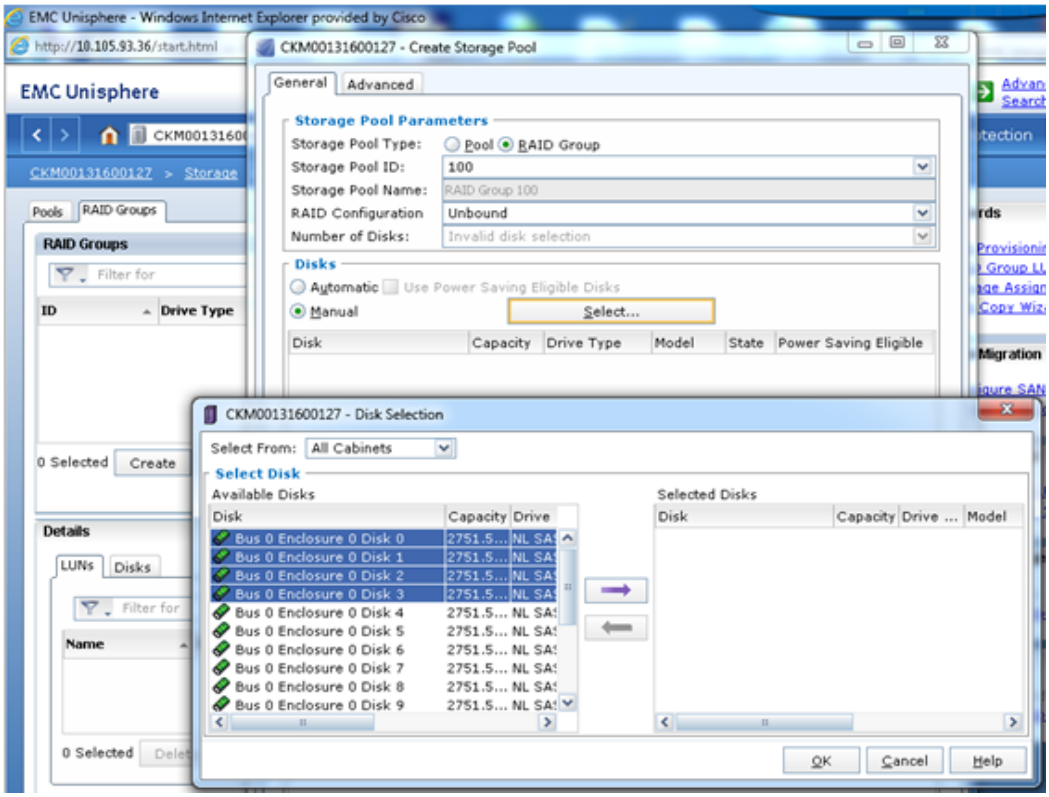
Copyright © 2014, Cisco Systems, Inc.

1. Go to **Storage** > **Storage Pools** > **RAID Group** and click **Create** as shown below:

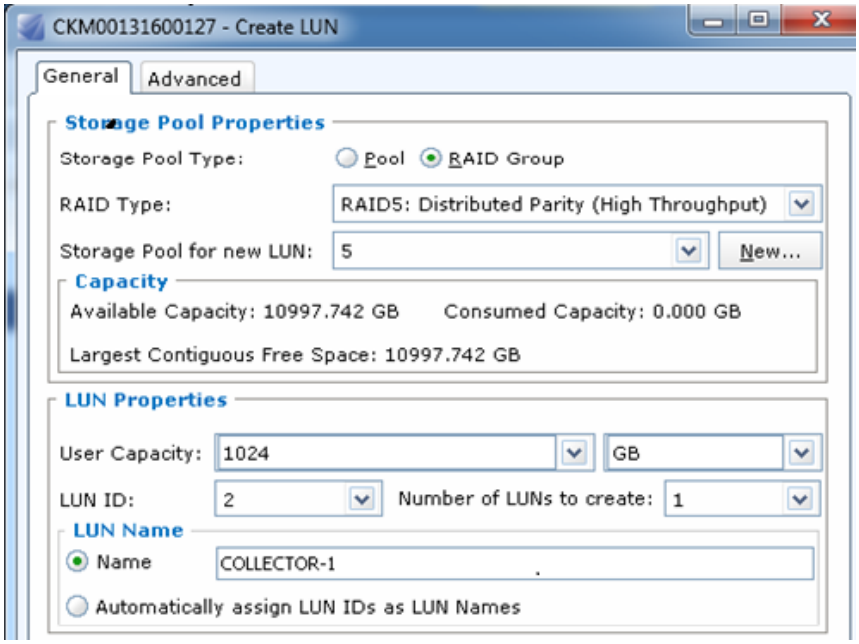


2. Select **RAID Configuration**> **Manual**.

3. Select disks, as planned earlier in this part of setting up the EMC by moving them from the **Available** list to the **Selected** list.



4. To create LUNs, go to **Storage** > **LUNS** > **LUNS** and click **Create**. The following screen displays:



5. In this screen, refer to the LUN creation table above, and specify the following parameters:

   - Select **Storage Pool Type** of RAID group.
   - Select a Storage Pool for the new LUN.
   - Select to match the RAID level. Note that the Insta nodes should be RAID type 10, while the other nodes should be RAID type 5.
   - Under **LUN Properties** > **User Capacity**, select the closest value (in GB) to that shown in the table above.
   - Select the LUN ID from the drop-down menu. LUN IDs auto-increment. Disks that are already assigned are not available.
   - Enter the LUN Name.

6. Click **Apply**.

7. Repeat Steps 4 through 6 for each of the nodes listed in the table above.

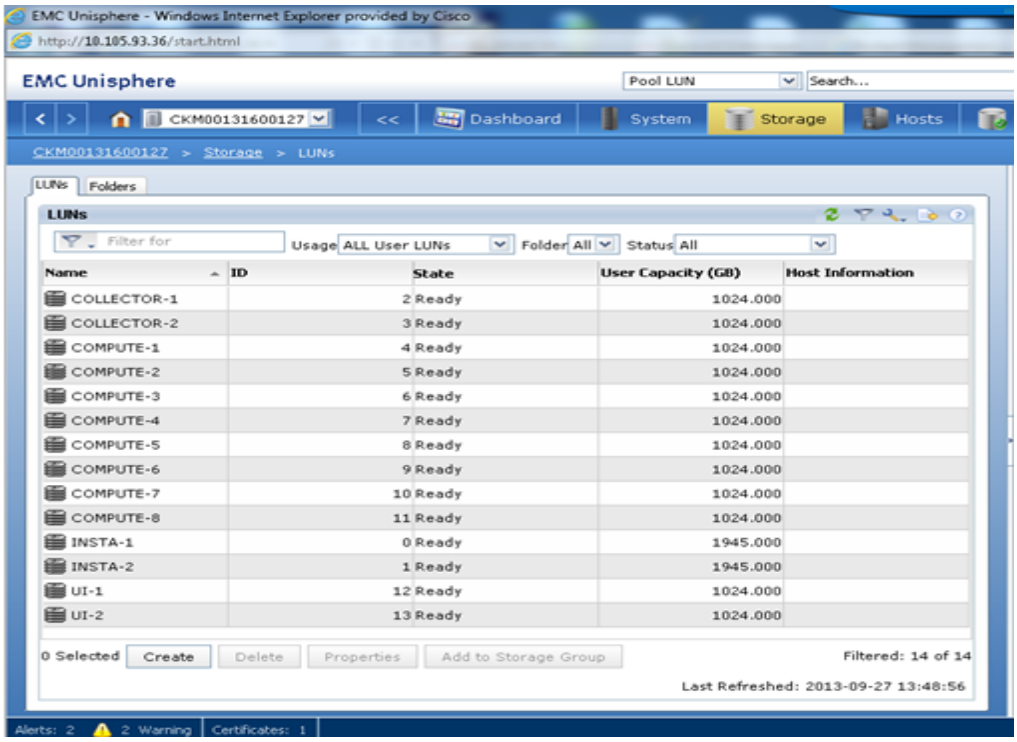
   Note that for the Insta nodes, you must create a total of four LUNs, each in RAID Group 10 (while the other types of nodes are in RAID Group 5):

Copyright © 2014, Cisco Systems, Inc.

- INSTA-1 (to be associated with dbroot1 on both Insta nodes)
- INSTA-2 (to be associated with dbroot2 on both Insta nodes)
- INSTA1-PGSQL (to be associated with INSTA-1 only)
- INSTA2-PGSQL (to be associated with INSTA-2 only)

The PGSQL storage partitions are for the Postgres database and will store Rubix-related data.

8. Verify that all LUNs have been correctly created and that their parameters match the LUN creation table above, by going to **Storage** > **LUNS**. A screen similar to the following displays:



Creating Storage Groups

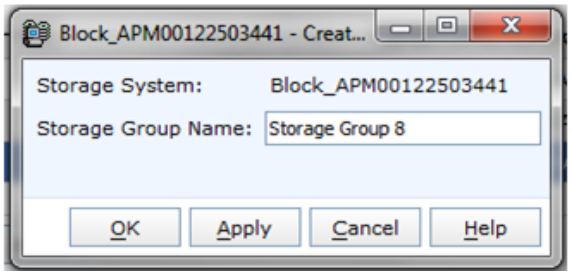
For Collector, Compute, Rubix, and GMS nodes, assigning the nodes to storage groups is a straightforward one-to-one relationship—the LUN assigned to each node has its own storage group.

However, the Insta nodes require two storage groups, each of which contains three LUNs:

- **INSTA1-PGSQL Storage Group 14**
  - INSTA-1
  - INSTA-2
  - INSTA1-PGSQL
- **INSTA1-PGSQL Storage Group 15**
  - INSTA-1
  - INSTA-2
  - INSTA2-PGSQL

To create the required storage groups for the MURAL system and assign LUNs to them:

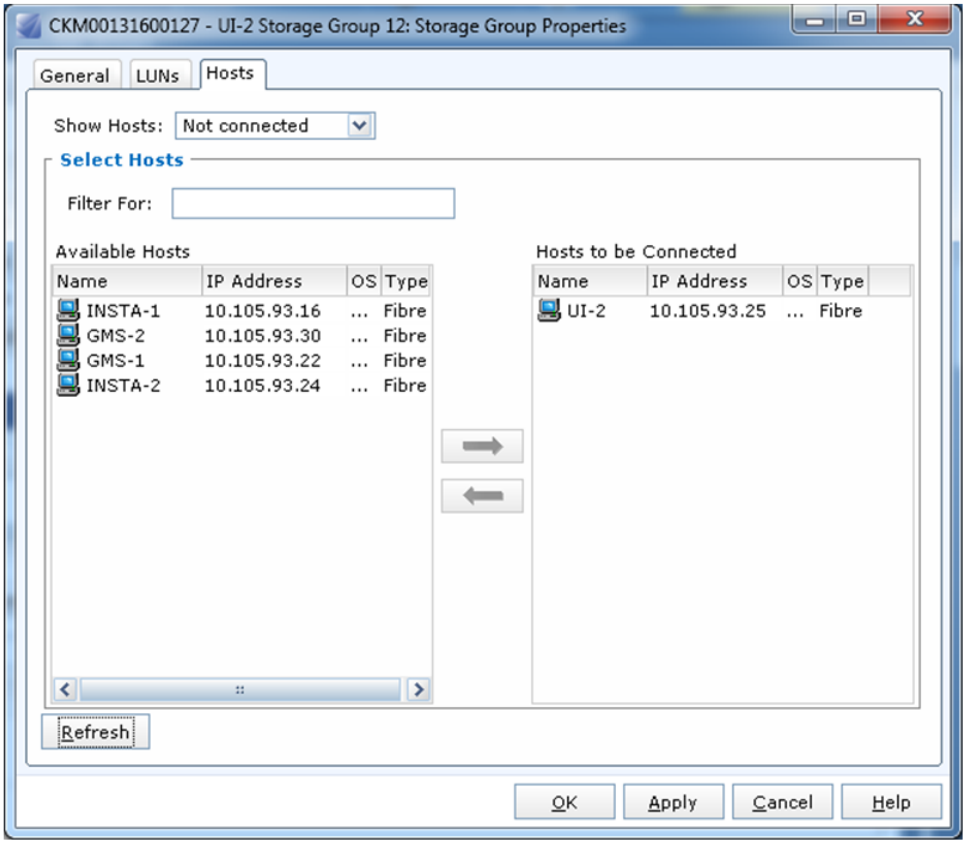
1. Go to **Hosts** > **Storage Groups**, and click **Create**. The following screen displays:



   The storage group number auto-increments.

Copyright © 2014, Cisco Systems, Inc.

2. The **Properties** dialogue box opens. From the **Hosts** tab, move the appropriate node to the **Hosts to be Connected** column:



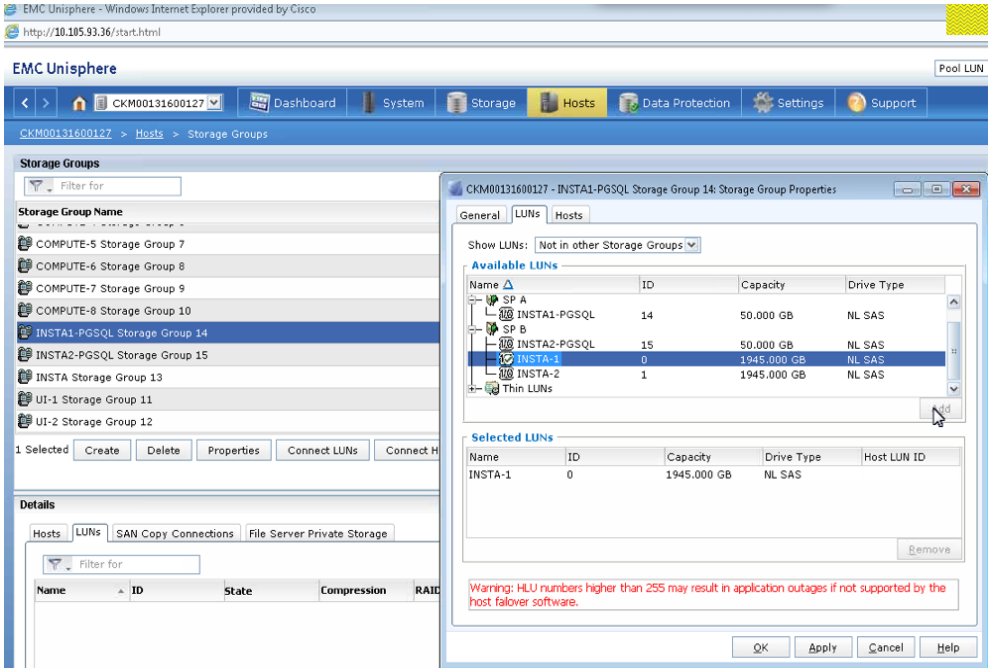
3. Click **OK**, then **Apply**.

4. Repeat Steps 1through 3 to create storage groups and assign the associated host name for each of the following nodes:

   - GMS 1
   - GMS 2 (if present in your environment)
   - Collector 1
   - Collector 2
   - Compute (also called DN for Data Node) 1
   - Compute 2
   - UI 1
   - UI 2

5. Repeat Steps 1-3 to create an Insta1-PGSQL storage group and from the **Hosts** tab, move host **INSTA-1** to the **Hosts to be connected** column.

6. Go to the **LUNs** tab in the properties window for the Insta1-PGSQL storage group and select LUN **INSTA-1**, and click **Add**.



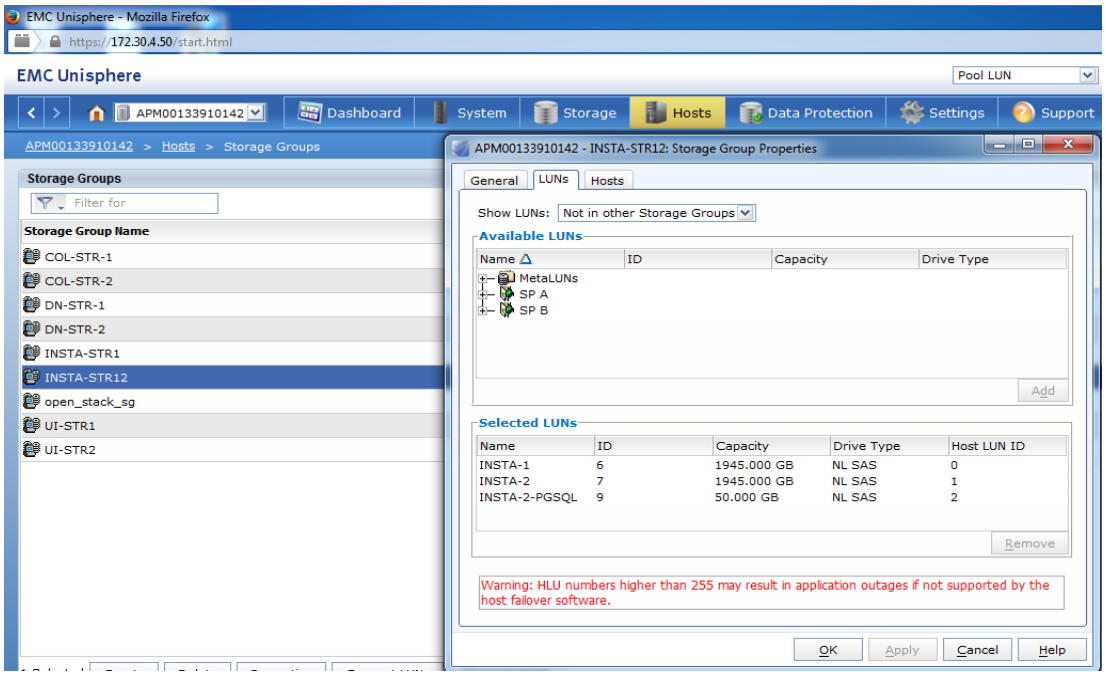
Then select LUN **INSTA-2** and click **Add**, and LUN **INSTA-1-PGSQL** and click **Add**.

7. Click **Apply**, and click **Yes** to confirm.

8. Repeat Steps 1 through 7 to create an Insta2-PGSQL storage group containing the LUNs **INSTA-1**, **INSTA-2**, and **INSTA2-PGSQL**.

9. Verify that all storage groups are listed similar to the example shown in the following screen (which does not show the required GMS storage group):

Copyright © 2014, Cisco Systems, Inc.



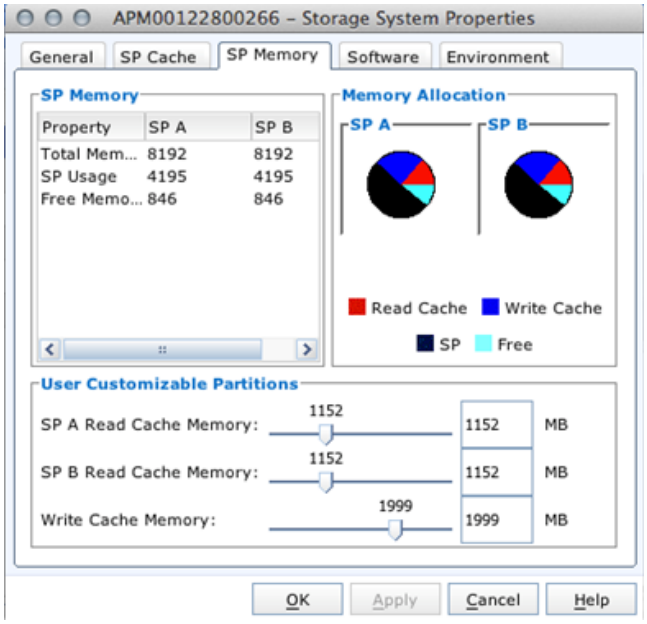
Adjusting Caching on the EMC

You must adjust the caching settings on the EMC for your MURAL setup to keep data in memory for the correct amount of time.

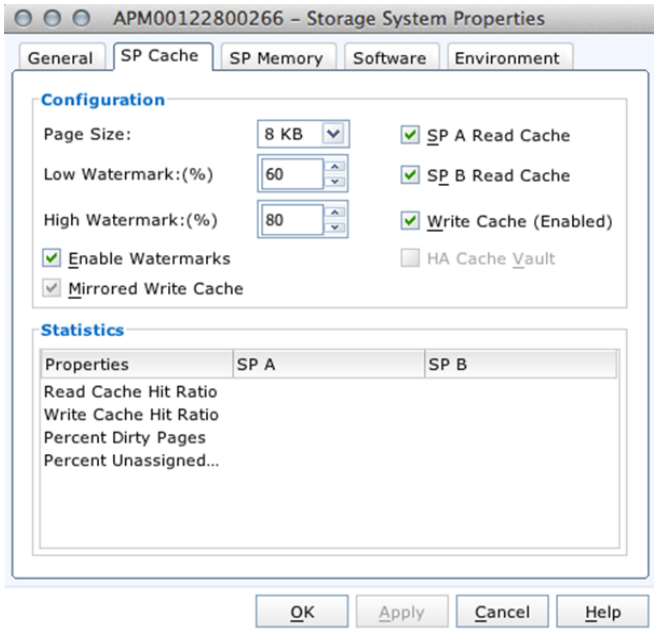
To adjust caching:

1. Select **System Management** > **Manage Cache**.
2. Disable all caching under **SP Cache**.

3. Adjust the SP memory, as shown in the example screen below: maximum to write and read set to 1152.



4. After completing the changes, re-enable all cache check boxes under the **SP Cache** tab.



EMC is now configured for MURAL.

Copyright © 2014, Cisco Systems, Inc.

Return to GMS and Configure the LUN WWIDs

Now that you have configured both UCS and EMC, return to the GMS to configure the actual WWIDs for the LUNs used in the MURAL system. During the initial configuration you used dummy values as placeholders, and now that the actual WWIDs are available, you must specify them in the GMS.

For instructions on configuring the WWIDs, see "Defining the Deployment Topology Using GMS" on page 55.

Installing MURAL on the UCS Nodes

In previous steps, you set the configuration parameters for your environment in the **mural.xml** file. Now you need to load this file onto each blade.

Before You Begin

Complete these tasks before you install MURAL on the UCS nodes.

- Update WWIDs on GMS configuration

  Before proceeding to the next section of this guide, ensure that the GMS configuration has been updated with WWID of all LUNs. See "Defining the Deployment Topology Using GMS" on page 55. During the initial configuration of the deployment topology, you used dummy values for the WWIDs of the LUNs. Make sure that these have been replaced with the actual values.

- Apply all available patches on the appropriate nodes.

  Before beginning the next section of this installation guide, patches and configuration changes must be applied on the nodes. To apply patches, follow the instructions in the Release Notes for software version 3.2 MR2 or above.

To install MURAL on the UCS nodes:

1. SSH to the GMS server using the management IP address and start the installation on all UCS nodes (fetches the software and puts it on the blade):

   ```
   > en
   # conf t
   (config) # install appliance all force-format
   ```

2. Monitor the installation status on all UCS blades:

   ```
   (config) # install appliance show installation-status all
   ```

   Note: The installation process takes approximately one hour to complete, although it may vary by installation site.

 Copyright © 2014, Cisco Systems, Inc.

The above command shows the percentage of the installation status per blade server. When the installation on all nodes is complete, the following messages are displayed:

```
     Collector-GMS-1 : Node successfully installed
Collector-GMS-2 : Node successfully installed
Compute-GMS-1 : Node successfully installed
CachingCompute-GMS-2 : Node successfully installed
Rubix-GMS-1 : Node successfully installed
Rubix-GMS-2 : Node successfully installed
```

Troubleshooting Node Installation

If the installation fails at any point, please contact Technical Support.

Logs can be collected on the GMS server from the location **/data/gms/logs**.

Tuning the System for Performance

To tune the performance of the system:

1. Log in to the master Collector node and run the following commands:

```
> en
# configure terminal
(config)# internal set modify -
/tps/process/hadoop/attribute/mapred.min.split.size/
value value string 268435456
(config)# internal set create - /tps/process/hadoop/attribute/
mapred.tasktracker.map.tasks.maximum/value value string 10
(config) # pm process tps restart
# write memory
```

2. Repeat Step 1 on the standby Collector node.

Verify the Processes are Running

After making any performance related modifications, verify that the processes are running.

Note: After you make any modifications, wait at least 15 minutes to ensure processes have had a chance to restart before running the following commands.

1. Log in to the master Collector node and verify that all Compute nodes have joined the HDFS cluster. The output shown in the following sample command sequence indicates correct configuration.

```
> en
(config) # _shell
 # ps -ef | grep hadoop | awk {'print $NF'}
 org.apache.hadoop.hdfs.server.datanode.DataNode
 org.apache.hadoop.hdfs.server.namenode.NameNode
 hadoop
 org.apache.hadoop.hdfs.server.namenode.NameNode
 org.apache.hadoop.hdfs.server.namenode.SecondaryNameNode
 org.apache.hadoop.mapred.JobTracker
# hadoop dfsadmin -report 2>/dev/null | egrep
"available|Name|Status"
 Datanodes available: 3 (3 total, 0 dead)
 Name: 10.10.2.13:50010
 Decommission Status : Normal
 Name: 10.10.2.14:50010
 Decommission Status : Normal
 Name: 10.10.2.17:50010
 Decommission Status : Normal
```

Note: These are internal IP addresses

```
# cli -t "en" "conf t" "show pm process collector" | grep status
      Current status:  running
```

2. Run the following commands on the standby Collector node:

Copyright © 2014, Cisco Systems, Inc.

```
> en
(config) # _shell
 #  ps -ef | grep hadoop | awk {'print $NF'}
 org.apache.hadoop.hdfs.server.datanode.DataNode
 org.apache.hadoop.hdfs.server.namenode.NameNode
 hadoop

# cli -t "en" "conf t" "show pm process collector" | grep status
      Current status:  running
```

3. Log in to the master Insta node, and run the following commands:

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

If the process is running, the command output is `Current status: running`.

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

The command output should show all modules in `ACTIVE` state.

It should also list all instances and Adaptor status as `RUNNING`.

4. Log in to the standby Insta node and repeat the same commands shown in step 3.

5. Run the following command to verify that Postgres is running on the master Insta node:

```
> en
# _shell
# ps -ef |grep pgsql
postgres  2990     1  0 Sep04 ?    00:00:01 /usr/pgsql-
9.2/bin/postmaster -p 5432 -D /data/pgsql/9.2/data
```

6. Repeat step 5 on the standby Insta node.

Generate and Push the Information Bases

To configure your system for the data specific for your environment, you must update the information bases (IBs) and fetch them. In some cases, you might need to also manually modify the IBs to add or reclassify entries.

Tethering

MURAL is shipped with initial basic Type Allocation Code (TAC) database for devices. This list aids the ASR in detecting tethered traffic. During initial deployment, you might want to add TACs to the list.

The following is the list of valid device groups returned from the Traffic Cat engine, depending on its TAC.

- 3G Watch
- Feature Phone
- M2M
- Mobile Router
- Netbook
- PC Card
- SmartPhone
- Tablet
- UNKNOWN

IB Modifications for the Tethering Feature

In addition to the necessary configurations done by the scripts, you might need to do some manual modifications of IBs in order to configure tethering for your environment, for example adding device groups for smart devices that you want to be considered in detection of tethering.

You must also create a file (**/data/work/serverFile_tethering**) to contain the details of the ASR gateways. The TAC, OS, and UA databases (the three reports associated with tethering) are created and stored in this file. The **serverFile_ tethering** file contains the entries for the data transfer destination location.

Note: The delimiter in this file must be ", " (comma followed by a space).

 Copyright © 2014, Cisco Systems, Inc.

For the SCP protocol, the destination folder should be present on the destination server. This is not required for SFTP.

To add groups to **smartDG.list** that you want to be considered as smart devices for tethering:

1. Determine which additional entries you might need to add by viewing device groups list on the master Collector node:

```
# conf t
 (config) # pmx
 pm extension > subshell bulkstats
 pm extension (bulkstats)> update all ibs from image
 pm extension (bulkstats)> quit
 pm extension > subshell aggregation_center
 pm extension (aggregation center) > update all ibs from image

 pm extension (aggregation center) > show ib deviceGroup.list
```

2. Choose from the names listed in the previous step and add them to the **smartDG** list:

```
pm extension (aggregation center)> edit ib smartDG.list add
Device Group: SmartPhone
pm extension (aggregation center)> edit ib smartDG.list add
Device Group: Feature Phone
pm extension (aggregation center)> show ib smartDG.list
       1 SmartPhone
       2 Feature Phone
pm extension (aggregation center)>
```

3. Create a file (**/data/work/serverFile_tethering**) to contain the details of the ASR gateways.

   To create the file, log in to the master Collector node:

```
> en
# _shell
```

```
# cd /data
# mkdir work
# cd work
# vi /data/work/serverFile_tethering
192.168.156.96, admin, admin, /data/TTYT
```

Include an entry for each ASR gateway. The file has the following format:

IP-address-of-gateway, *ASR-gateway-user-login-name*, *ASR-gateway-password*, *location-on-ASR-gateway-for-databases-copies*

For example:

```
192.168.1.1, admin, password, /data
```

4. Repeat the previous step on the standby Collector node.

Configuring IBs for EDR

The following table shows a sample data set for setting up the IBs.

Sample Data Set

DC	Gurgaon	GGSNIP	27.23.157.1	SGSNIP	2.2.2.1	SGSN	CYBERCITY	APN	AIRTEL-NH8	GROUP	AIRTEL	RATID	1	RATTYPE	CDMA
Region	SEZ	GGSN	GURGAON-GGSN	SGSNIP	2.2.2.2	SGSN	UBICITY	APN	AIRTEL-GGN	GROUP	AIRTEL	RATID	2	RATTYPE	GSM
Area	ITPARKS							APN	VDF-NH8	GROUP	VODAFONE	RATID	3	RATTYPE	WIMAX
								APN	Sushfone-1	GROUP	VODAFONE	RATID	4	RATTYPE	LTE

Note: Use the above table for example purposes only. You should use the data that matches your environment. For example, for GGSN, you might use GGSN, PGW, or HA. In this case, GGSNIP is the management IP address. For SGSN, you might use SGSN, SGW, HSGW, or PDSN. In this case, SSGNIP is the service IP address.

To configure the IBs for EDR:

Tip: Repeat the `add` command to add the IP address of each node in turn.

1. Login to the master Collector node, and add Control or Internal IP(s) addresses of all Collector and UI nodes; that is, include all master and standby nodes of both collector and UI nodes:

```
> en
```

 Copyright © 2014, Cisco Systems, Inc.

```
# conf t
(config) # pmx
 pm extension> subshell aggregation_center
 pm extension (aggregation center)> add ib-destination-IP
 Control or Internal IP  addresses of all master and all
 /standby collectors and UIs nodes
```

2. Enter the GGSN IPs and GGSN names:

```
pm extension (aggregation center)> edit ib ipGgsn.map add
```

3. Enter the SGSN IPs and SGSN name:

```
pm extension (aggregation center)> edit ib ipSgsn.map add
```

4. Enter the APN name and corresponding APN group:

```
pm extension (aggregation center)> edit ib apnGroup.map add
```

5. Enter the RAT ID and RAT TYPE:

```
pm extension (aggregation center)> edit ib ratidtype.map add
```

6. Change the default segments to have only **Low**, **Regular**, and **High** segments in the configuration, as shown in the following example:

```
pm extension> subshell aggregation_center `
pm extension (aggregation center)> edit ib segment.map delete
record 5
IB [segment.map] updated.
pm extension (aggregation center)> show ib segment.map
     1  [629145600][Extreme]
     2  [314572800][High]
     3  [157286400][Active]
     4  [62914560] [Regular]
     5  [0]        [Low]
pm extension (aggregation center)> edit ib segment.map delete
record 3
IB [segment.map] updated.
pm extension (aggregation center)> show ib segment.map
     1  [629145600][Extreme]
     2  [314572800][High]
     3  [62914560] [Regular]
```

```
     4  [0]         [Low]
pm extension (aggregation center)> edit ib segment.map delete
record 1
IB [segment.map] updated.
pm extension (aggregation center)> show ib segment.map
     1  [314572800][High]
     2  [62914560] [Regular]
     3  [0]         [Low]
```

Configuring DCs and Gateways For All IBs

Add all new DC/gateways to the configuration for the system. DC, ASR, and Gateway are synonymous terms and hold same meaning for the system. Gateway/DC name is a unique key for a gateway configuration. In the BulkStat hierarchy, the gateway name and DC name are the same.

Guidelines for adding gateways:

- All input directories are created under the **/data/collector** path. Hence, in the example below, the ASR should send EDR files to **/data/collector/California/edr111** and send the Bulkstats file to **/data/collector/California/bs111**.

  Important: Confirm the actual input directories with the Cisco Technical Support team for your ASR platform.
- Ensure that you provide different input paths for each new gateway being added.
- The ***/edr-file-path*** and ***/bulkstat-file-path*** should always start with a forward slash (/).
- Files coming for BulkStats should contain the GatewayName delimited by an underscore (_). For example, GMPLAB1_20130912121500.
- The filename pattern provided here should be in sync with the Collector configurations.
- The ASR should send the gateway name in place of %*DC*, as specified in the file name pattern in the Collector configurations.
- If the filenames will have file extensions of **.gz** or **.txt**, then you must provide ".*" in the file name format configuration when adding the gateway.

 Copyright © 2014, Cisco Systems, Inc.

- All incoming files should contain the string as per their type in the file name; that is, flow EDR files should contain the string "flow," delimited by an underscore (_) and HTTP EDR files must contain string "http," delimited by an underscore (_).

To configure gateways:

1. Add gateway information for each gateway (refer to guidelines above for clarification about input parameters):

```
> en
# conf t
(config)# pmx subshell aggregation_center
pm extension (aggregation center)> add gateway namegateway-name
region gateway-region location gateway-area schema_version
bulkstat-schema-version ip gateway-IP timezone gateway-timezone
edr-filename-pattern incoming-EDR-fileName-pattern
bulkstat-filename-pattern incoming-BS-fileName-pattern
type gateway-type edr-file-path incoming-EDR-files-path-on-collector
bulkstat-file-path incoming-BS-files-path-on-collector
```

   Tip: For `bulkstat-schema-version`, specify aninteger only, for example use `15`, do not use `15.0`.

   For example:

```
pm extension (aggregation center)> add gateway name GMPLAB1
region EAST location USA
schema_version 15 ip 10.10.10.255
timezone America/Recife
edr-filename-pattern *_MURAL-edr_*_%MM%DD%YYYY%hh%mm%ss.* bulkstat-
filename-pattern
bulkstats_%YYYY%MM%DD%hh%mm%ss type HA edr-file-path /California/edr111
bulkstat-file-path /California/bs111
```

   Example output:

```
Adding IBs....
*************************Adding ib
dcRegionArea*****************************
Adding in dcRegionArea.map
```

```
[dcRegionArea.map]:
    generated dc.id
    generated region.id
    generated area.id
    generated dcRegionArea.id.map

Summary:
========
Successful IBs : 1 out of 1
Failed IBs : No id generation failures.
pushing ib [dc.id] to all IB destinations
pushing ib to ip 192.168.151.173
This product may be covered by one or more U.S. or foreign pending or
issued patents listed at www.guavus.com/patents.

Guavus Network Systems NetReflex Platform

dc.id

             100%  133     0.1KB/s   00:00
This product may be covered by one or more U.S. or foreign pending or
issued patents listed at www.guavus.com/patents.

Summary:
========
Successful Transfers : 4 out of 4
Failed Transfers : No transfer failures.
pushing ib [region.id] to all IB destinations
pushing ib to ip 192.168.151.173
This product may be covered by one or more U.S. or foreign pending or
issued patents listed at www.guavus.com/patents.

Guavus Network Systems NetReflex Platform

region.id

             100%   86     0.1KB/s   00:00
```

 Copyright © 2014, Cisco Systems, Inc.

```
**************************Added ib
dcRegionArea*********************************

**************************Adding ib
gateway************************************
Adding in gateway.map
[key.map]:
    generated key.id.map
[gateway.map]:
    generated gateway_bs.id
    generated version_bs.id
    generated dc_bs.id
    generated region_bs.id
    generated area_bs.id
    generated gatewayRegionArea.id.map
[schema.map]:
    generated schema.id.map
[metric.map]:
    generated metric.id.map
[gatewayIDVersion.map]:

Summary:
========
Successful IBs : 5 out of 5
Failed IBs : No id generation failures.
pushing all ibs to all IB destinations
pushing ib to ip 192.168.151.173
This product may be covered by one or more U.S. or foreign pending or
issued patents listed at www.guavus.com/patents.

Guavus Network Systems NetReflex Platform

gatewayIDVersion.map 100%  113     0.1KB/s   00:00
schemaIDToKeyID.map 100% 8018     7.8KB/s   00:00
Gateway_Schema_ASR.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_GMPLAB1.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_GMPLAB2.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_NewDelhi.map 100% 1779KB   1.7MB/s   00:00
```

```
Gateway_Schema_gmplab1.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_gmplab2.map 100% 1779KB   1.7MB/s   00:00
gatewaySchemaMetric.map 100%   10MB  10.4MB/s   00:00
key.id.map 100% 1375     1.3KB/s   00:00
key.map 100%  748     0.7KB/s   00:00
gateway_bs.id 100%  113     0.1KB/s   00:00
version_bs.id 100%   12     0.0KB/s   00:00
dc_bs.id 100%  113     0.1KB/s   00:00
region_bs.id 100%   31     0.0KB/s   00:00
area_bs.id 100%   34     0.0KB/s   00:00
gatewayRegionArea.id.map 100%  228     0.2KB/s   00:00
gateway.map 100%  178     0.2KB/s   00:00
schema.id.map 100% 1883     1.8KB/s   00:00
schema.map 100%  816     0.8KB/s   00:00
subschema_definition.list 100% 1007KB   1.0MB/s   00:00
metric.id.map 100%  679KB 679.1KB/s   00:00
metric.map 100%  471KB 471.1KB/s   00:00
cube_names.list 100%  384     0.4KB/s   00:00
schemaIDToKeyID.map 100% 8018     7.8KB/s   00:00
Gateway_Schema_ASR.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_GMPLAB1.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_GMPLAB2.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_NewDelhi.map 100% 1779KB   1.7MB/s   00:00
Gateway_Schema_gmplab1.map 100% 1779KB   1.7MB/s   00:01
Gateway_Schema_gmplab2.map 100% 1779KB   1.7MB/s   00:00
gatewaySchemaMetric.map 100%   10MB  10.4MB/s   00:00

This product may be covered by one or more U.S. or foreign pending or
issued patents listed at www.guavus.com/patents.

Guavus Network Systems NetReflex Platform

_DONE

             100%    0     0.0KB/s   00:00

Summary:
========
```

 Copyright © 2014, Cisco Systems, Inc.

```
Successful Transfers : 4 out of 4
Failed Transfers : No transfer failures.
**************************Added ib
gateway***********************************
Adding Gateway configs....
************************Gateway*********************************
Name: gmplab2
Associated Region: EAST
Location: USA
Schema Version: 14
IP: 10.10.10.255
Timezone: America/Recife
Flow-EDR/Http-EDR Filename Pattern: %DC_MURAL-edr_
*_%MM%DD%YYYY%hh%mm%ss*.gz
Bulkstat Filename Pattern: *_%YYYY%MM%DD%hh%mm%ss
Type: HA
Flow-EDR/Http-EDR Filename Path: /data/collector/California/edr111
Bulkstat Filename Path: /data/collector/Graham/bs111
*********************Successfully Added************************
```

2. Verify the new gateway has been added:

   ```
   pm extension (aggregation center)> show gateways
   ```

3. Add IPs for all Collectors to push this information to all Collectors:

   ```
   pm extension (aggregation center)> set collector IPs comma-separated-ip-list
   ```

   Example:
   pm extension (aggregation center)> **set collector IPs 192.168.1.1, 192.168.2.2**

   Note: These are internal IP addresses.

4. Verify the IP addresses for all Collectors:

   ```
   pm extension (aggregation center)> show collector IPs
   ```

5. Set the BulkStats timezone to **UTC** in **gateway.conf** for every gateway. The reason for this is that the ASR internally changes the time zone to GMT

for the BulkStats file. Edit **gateway.conf** for every BulkStats source at the path:

```
/data/configs/gateway/gateway.conf "timezone": "UTC"
```

6. Push the gateway configuration to all the Collectors:

```
pm extension (aggregation center)> push gateway configuration
```

7. Generate and push all IBs:

```
pm extension (aggregation center)> generate all ibs
pm extension (aggregation center)> push all ibs
pm extension (aggregation center)> quit

pm extension> quit
(config) # write memory
```

Synchronize the EDR and BulkStats IBs on the Standby Collector Node

After the above **push** command completes, run the following command on the standby Collector node from the CLI configure terminal.

To fetch IBs for EDR data streams:

```
pmx
 pm extension> subshell aggregation_center
 pm extension (aggregation center)> fetch all ibs from inbox
 Copying IB mobileappname.list [OK]
 Copying IB segment.map [OK]
 Copying IB apnGroup.map [OK]
 Copying IB snAppCat.map [OK]
 Copying IB privateKey.txt [OK]
 Copying IB apnNetworkIB.map [OK]
 Copying IB segment.txt [OK]
 Copying IB dcNameClli.map [OK]
 Copying IB ReservedNames.json [OK]
 Copying IB collection_center.list [OK]
 Copying IB ratidtype.map [OK]
 Copying IB wngib.id.map [OK]
```

 Copyright © 2014, Cisco Systems, Inc.

```
Copying IB pcsarange.list [OK]
Copying IB ServiceGateway.list [OK]
Copying IB dcRegionArea.map [OK]
Copying IB mobileappcategory.list [OK]
Copying IB model.list [OK]
Copying IB snAppIdStr.map [OK]
Copying IB UnresolvedTraps [OK]
Copying IB subdevice.id.mph [OK]
Copying IB ipGgsn.map [OK]
Copying IB keyword.list [OK]
Copying IB sp.list [OK]
Copying IB url.list [OK]
Copying IB ipSgsn.map [OK]
Copying IB SNMPRecordConverterConfig [OK]
Copying IB blacklist [OK]
Copying IB mobileapptype.list [OK]
Copying IB ivFile.txt [OK]
Copying IB smartDG.list [OK]
Copying IB ipProto.map [OK]
Copying IB mime.list [OK]
Copying IB topsubscriberib.id.map [OK]
Copying IB p2pCat.map [OK]
Copying IB tac_manuf_model.map [OK]
Copying IB subseg.id.map [OK]
Copying IB sp_url_cat.map [OK]
Copying IB symmetricKey.txt [OK]
Copying IB manufacturer.list [OK]
Copying IB IBStore.tab [OK]
Copying IB category.list [OK]
Copying IB ua_mobile_app.map [OK]
Copying IB subscriberib.id.map [OK]
Copying IB port_proto_app.map [OK]
Copying IB ua_manuf_model.map [OK]
Copying IB mobileappname.id [OK]
Copying IB bytes.id [OK]
Copying IB segmentname.id [OK]
Copying IB segment.id.map [OK]
Copying IB apn.id [OK]
```

```
Copying IB group.id [OK]
Copying IB apnGroup.id.map [OK]
Copying IB snAppProtocolStr.id [OK]
Copying IB snAppCategory.id [OK]
Copying IB snAppCat.id.map [OK]
Copying IB ratid.id [OK]
Copying IB rattype.id [OK]
Copying IB ratidtype.id.map [OK]
Copying IB dc.id [OK]
Copying IB region.id [OK]
Copying IB area.id [OK]
Copying IB dcRegionArea.id.map [OK]
Copying IB mobileappcategory.id [OK]
Copying IB model.id [OK]
Copying IB snAppid.id [OK]
Copying IB snAppStr.id [OK]
Copying IB snAppIdStr.id.map [OK]
Copying IB ggsnip.id [OK]
Copying IB ggsn.id [OK]
Copying IB ipGgsn.id.map [OK]
Copying IB sp.id [OK]
Copying IB url.id [OK]
Copying IB sgsnip.id [OK]
Copying IB sgsn.id [OK]
Copying IB ipSgsn.id.map [OK]
Copying IB mobileapptype.id [OK]
Copying IB ipProtocolId.id [OK]
Copying IB ipProtocolStr.id [OK]
Copying IB ipProto.id.map [OK]
Copying IB mime.id [OK]
Copying IB p2pProtocolStr.id [OK]
Copying IB p2pCategory.id [OK]
Copying IB p2pCat.id.map [OK]
Copying IB manufacturer.id [OK]
Copying IB category.id [OK]
Copying IB Destport.id [OK]
Copying IB ipProtoStr.id [OK]
Copying IB appProtoStr.id [OK]
```

 Copyright © 2014, Cisco Systems, Inc.

```
Copying IB port_proto_app.id.map [OK]
pm extension (aggregation center)>
pm extension (aggregation center)> quit
pm extension> quit
COL-02 [COL-VIP: standby] (config) # write mem
COL-02 [COL-VIP: standby] (config) #
```

Fetching BulkStats IBs

To configure your MURAL environment to evaluate BulkStats data, go to the bulk-stats PMX subshell and fetch the IBs from the inbox:

```
(config) # pmx
pm extension> subshell bulkstats
pm extension (bulk stats)> fetch all ibs from inbox
Copying IB gatewaySchemaMetric.map            [OK]
Copying IB gatewayIDVersion.map               [OK]
Copying IB key.map                            [OK]
Copying IB gateway.map                        [OK]
Copying IB schema.map                         [OK]
Copying IB subschema_definition.list          [OK]
Copying IB metric.map                         [OK]
Copying IB cube_names.list                    [OK]
Copying IB key.id.map                         [OK]
Copying IB gateway_bs.id                      [OK]
Copying IB version_bs.id                      [OK]
Copying IB dc_bs.id                           [OK]
Copying IB region_bs.id                       [OK]
Copying IB area_bs.id                         [OK]
Copying IB gatewayRegionArea.id.map           [OK]
Copying IB schema.id.map                      [OK]
Copying IB metric.id.map                      [OK]
pm extension (bulk stats)>
```

Create File for Uncategorized URL, UA, and TAC Reports

To create the file needed for reports for uncategorized URLs, UAs, and TACs:

1. Create a file on the master Collector node containing the destination information, to which the uncategorized URL, UA, and TAC reports will be

copied, as shown:

File : /data/work/serverFile_uncatReports

"serverFile_uncatReports" contains the entry for data transfer destination. This file has the following format:

IP, *username*, *password*, *location-to-which-reports-should-be-copied*

Example:

```
192.168.156.96, admin, password, /data/offline_uncat_reports
```

Note: The delimiter in this file must be ", " (comma followed by a space).

2. Log in to the master Collector node:

```
> en
# _shell
# cd /data
# cd work
# vi /data/work/serverFile_uncatReports

192.168.156.96, admin, password, /data/TTYT
```

Note: For the SCP protocol, the destination path should be present at the destination server. The destination path is not required for SFTP.

3. Create the same file on the standby Collector node.

 Copyright © 2014, Cisco Systems, Inc.

Single Certificate Installation

This section provides instructions for generating and installing a common certificate for all Tomcat processes on the distributed Rubix setup. The example shown is based on generating and using a certificate with the Microsoft CA. However, when deploying the MURAL application, you can use any third party CA (Verisign, for example).

Note: MURAL supports the Mozilla Firefox and Chrome browsers.

Backing Up and Generating the Keystore Files

1. Log in to the master Rubix node and make a backup of the original keystore.

   ```
   mv /data/apache-tomcat/apache-tomcat-7.0.27/keystore /data/apache-tomcat/apache-tomcat-7.0.27/keystore_orig
   ```

2. Generate the new keystore file following the steps below:

   ```
   cd /usr/java/latest/bin/
   ```

   ```
   ./keytool -keysize 2048 -genkey -alias tomcat -keyalg RSA -keystore keystore
   ```

   The keytool command prompts for the following values:

   - **password**—rubix123
   - **first and last name**—The fully-qualified domain name, or URL, you are securing. If you are requesting a wildcard certificate, add an asterisk (*) to the left of the common name where you want the wildcard; for example, ***.cisco.com.**
   - **Organizational Unit**—Optional
   - **Organization**—The full legal name of your organization (for example, Cisco)
   - **City/Locality**—Name of the city in which your organization is registered/located; do not abbreviate

- **State/Province**—Name of state or province where your organization is located; do not abbreviate
- **Country Code**—The two-letter International Organization for Standardization (ISO) format country code for where your organization is legally registered.

```
Enter keystore password:
Re-enter new password:
What is your first and last name?
  [Unknown]:  *.cisco.com
What is the name of your organizational unit?
  [Unknown]:
What is the name of your organization?
  [Unknown]:  Cisco
What is the name of your City or Locality?
  [Unknown]: Tewksbury
What is the name of your State or Province?
  [Unknown]: Massachusetts
What is the two-letter country code for this unit?
  [Unknown]:  US
Is CN=*.cisco.com, OU=Unknown, O=Guavus, L=Gurgaon,
ST=Haryana, C=IN correct?
  [no]:  yes
Enter key password for <tomcat>
    (RETURN if same as keystore password):
[admin@CISCOMUR-UI-147-17 bin]#
```

3. Run the following command to create a CSR file called **csr.csr**:

```
[admin@CISCOMUR-UI-147-17 bin]# ./keytool -certreq -keyalg RSA
-alias tomcat -file csr.csr -keystore keystore
```

When prompted for the keystore password, type the same password as for the last prompt in step 2.

4. Open the CSR file, and copy all of the text.

```
[admin@CISCOMUR-UI-147-17 bin]# vi csr.csr
```

 Copyright © 2014, Cisco Systems, Inc.

5. Paste all of the text into the online request form, as described in the following steps, and complete your application.

Downloading the Signed Certificate

1. Log in to **https://mx1.guavus.com/certsrv**.
2. Enter your username and password.
3. From the list of tasks, select **Request a certificate**.
4. Select **Advanced Certificate Request** from the next screen.
5. On the next screen, **Submit a Certificate Request or Renewal Request**, enter the CSR generated above (in the **csr.csr** file) and select **Web Server** from the Certificate Template.
6. Click **Download certificate** to download the CA certificate.

   The certificate you requested is issued.

7. Rename the downloaded file to **tomcat.cer**.
8. Place the file **tomcat.cer** on the same path on the server where the keystore is generated: **/usr/java/latest/bin/**

   ```
   scp tomcat.cer admin@172.16.10.0:/usr/java/latest/bin
   ```

Downloading the CA Certificate

1. Log in to **https://mx1.guavus.com/certsrv**.
2. Enter your username and password.
3. From the list of tasks, select the task **Download a CA certificate, certificate chain, or CRL**
4. On the next screen, select **Download CA certificate**.
5. Rename the downloaded file to **root.cer** and place the file on the same path on the server where the keystore is generated: **/usr/java/latest/bin/**

   ```
   scp root.cer admin@172.16.10.0:/usr/java/latest/bin
   ```

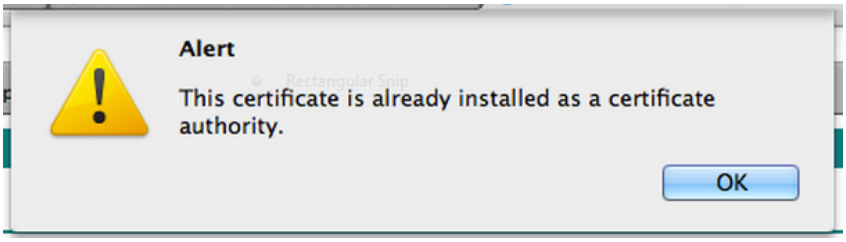
6. Click on the **Install CA certificate** link.

   Note: This option should be used on all client machines to install the CA certificate on the client, which will be able to identify the

certificates signed using GUAVUS-MX1-CA.

7. Click all the check boxes and then click **OK**.

   Note: If the CA certificate is already installed, the following message is displayed:



Installing the Signed Certificate in the Keystore

For the CA GUAVUS-MX1-CA, the root (CA) certificate and signed certificates (downloaded using the steps given above) need to be added to keystore.

1. Using keytool, enter the following commands to install the signed certificates and root certificate

   Enter the keystore password and type **yes** to accept the certificate:

```
[admin@CISCOMUR-UI-147-17 bin]# ./keytool -import -alias root -keystore
keystore -trustcacerts -file root.cer
Enter keystore password:
Owner: CN=Guavus-MX1-CA, DC=guavus, DC=com
Issuer: CN=Guavus-MX1-CA, DC=guavus, DC=com
Serial number: 7570961243691e8641088832d6c218a5
Valid from: Sun Feb 14 15:36:50 GMT 2010 until: Fri Feb 14 15:46:49 GMT
2020
Certificate fingerprints:
     MD5:  27:BE:E3:EA:DB:16:A1:79:8B:00:96:A4:54:D8:16:6D
     SHA1: 8B:6B:44:2B:08:FE:3F:C5:8B:B5:AD:FA:EE:DF:E2:A7:07:59:D0:49
     Signature algorithm name: SHA1withRSA
     Version: 3
Extensions:
#1: ObjectId: 2.5.29.19 Criticality=true
BasicConstraints:[
  CA:true
```


Copyright © 2014, Cisco Systems, Inc.

```
  PathLen:2147483647
]
#2: ObjectId: 2.5.29.15 Criticality=false
KeyUsage [
  DigitalSignature
  Key_CertSign
  Crl_Sign
]
#3: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: D4 DA C9 FC FB AF 5E 32   BF C0 A3 F1 6F 03 39 C3
......^2....o.9.
0010: 29 36 84 7D                                        )6..
]
]
#4: ObjectId: 1.3.6.1.4.1.311.21.1 Criticality=false
Trust this certificate? [no]:  yes
Certificate was added to keystore
```

2. Install the signed certificate:

```
[admin@CISCOMUR-UI-147-17 bin]# ./keytool -import -alias tomcat -
keystore keystore -trustcacerts -file tomcat.cer
```

3. Enter the keystore password.

   The certificate reply was installed in keystore

4. Copy the new keystore to the **apache-tomcat-7.0.27** directory for all EDRs, Bulkstats and RGE on both master and standby machines.

```
[admin@CISCOMUR-UI-147-17 bin]# cp keystore /data/apache-
tomcat/apache-tomcat-7.0.27/
[admin@CISCOUCS-2-19 apache-tomcat-7.0.27]# scp keystore
admin@192.168.147.20:/data/apache-tomcat/apache-tomcat-7.0.27/
```

   Note: 192.168.147.20 is the management IP address of the standby Rubix node.

5. Verify that the keystorePass in **server.xml** of all Tomcat processes is the same as the keystore password given above:

```
# vi /data/apache-tomcat/apache-tomcat-7.0.27/conf/server.xml
 <Connector port="8443" protocol="HTTP/1.1" SSLEnabled="true"
               maxThreads="150" scheme="https" secure="true"
               keystoreFile="keystore"
keystorePass="rubix123"
               clientAuth="false" ciphers=" SSL_RSA_WITH_RC4_128_MD5,
SSL_RSA_WITH_RC4_128_SHA" sslProtocol="TLS" />
```

 Copyright © 2014, Cisco Systems, Inc.

Processing the Data

This section includes information for setting up an ASR user in the Collector nodes, sending the EDR and BulkStats data feeds to the MURAL platform, setting the data start time, and running the data processing commands.

Use one of the ASR data feed methods to send data to the MURAL platform.

Setting Up a New User for the ASR in the Collectors

Note: The username and password should be the same ones configured on the ASR for EDR and BulkStats file transfers.

To set up a new user for the ASR in the Collectors:

1. Log on to the master Collector node and create the user:

```
# en
> conf t
(config)> username userid password password
(config)> write memory
(config)> _shell
```

2. Edit **/etc/ssh/sshd_config** to set the following parameters, in case a PAM related authentication error is reported in the master Collector node **/var/log/messages** for the ASR:

```
# mount -o remount,rw /
UsePAM no
PasswordAuthentication yes
```

3. Run the `pm process sshd restart` command:

```
# en
> conf t
(config) pm process sshd restart
(config) _shell
# mount -o remount,ro /
```

4. Repeat steps 1 through 3 on the standby Collector node.

Ingesting the Data Into the System

Start sending the EDR and BulkStats data feeds to the MURAL platform. If the ASR is used as an input node, the start time from the filename is created in the incoming EDR flow directory.

The file name has the timestamp, which can be used for job scheduling in the following process.

Note: It is assumed that the timestamp on the data that is pushed to the platform is greater than or equal to the current time, and not an old timestamp.

Copyright © 2014, Cisco Systems, Inc.

Validating the System Installations

This section includes instructions for validating data after completing the system installations.

Data Validation on the Collector Nodes

1. Login to the master Collector node:

   ```
   > en
   # _shell
   ```

2. Run the following two commands:

   ```
   # hadoop dfs -ls
   /data/collector/1/output/edrflow/YYYY/MM/DD/HH/mm/* 2>/dev/null
   # hadoop dfs -ls
   /data/collector/1/output/edrhttp/YYYY/MM/DD/HH/mm/* 2>/dev/null
   ```

   If the Collector is receiving data in the expected format, it will retain the data in HDFS.

   Note: Specify the year, month day, hour, and minute for which data is being sent to the MURAL system; minutes are always in multiples of 5 - 00,05,10,....55.

   These directories and files are updated continuously as the data keeps coming in.

3. Execute the following command:

   ```
   # hadoop dfs -ls
   /data/collector/1/output/bulkStats/YYYY/MM/DD/HH/mm/* 2>/dev/null
   ```

   If the Collector is receiving data in the expected format, it will retain the data in HDFS.

   Note: Specify the year, month, day, hour, and minute for which data is being sent to the MURAL system; minutes are always in multiples of 5 - 00,05,10,....55.

Set the Data Start Time

To set the data start time in the configuration:

1. Log in to the GMS server:

```
> en
# _shell
# mount -o remount,rw /
# cd /opt/deployment/Mural_setStartTime/
# ./setOozieTime --dataStartTime data_start_time --node
collector_mgmt_IP --password admin_user_pwd --verbose
```

2. Execute the script to set the data start times to the time from which EDR and BulkStats data starts coming into the Hadoop system from the previous section, "Data Validation on the Collector Node."

   For example, if EDR and Bulk Stats data starts coming into the Hadoop system from April 1, 2013, 06:00 onwards, run the following scripts with the start_time value as "2013-04-01T06:00Z":

```
# ./setOozieTime --dataStartTime  2013-04-01T06:00Z --node
192.168.147.11 --password admin@123 --verbose
```

   Note: Enter minutes as a multiple of 5. For example, "2013-04-01T06:00Z." Ensure that there is a continuous flow of data into the Hadoop without any gaps since the specified time.

3. Execute the Set Job Time Script for both the master and standby Collector nodes.

Start the Data Processing

1. Log in to the master Collector/Name node and run the data processing commands from the Oozie subshell:

```
> en
# conf t
(config)# pmx
Welcome to pmx configuration environment.
pm extension> subshell oozie
pm extension (oozie)> run job all
```

The command output shows all the jobs that were initiated and if the jobs started successfully or not.

 Copyright © 2014, Cisco Systems, Inc.

Data Validation on Compute Blades (Data Nodes)

This section includes the steps required to validate data on the Compute blades (Data nodes) following the installation process.

Note: Wait two hours after completing the above steps before doing the checks described in this section. This allows sufficient time for the jobs that process the collector data to start, and the **done.txt** files to be updated.

EDR Data

All commands in this section are to be executed on the master Collector node.

1. Check the last timestamp for EDR data cubes being generated by the EDR job from the master Collector node.

```
> en
# _shell
[admin@collector-1 0000041-121020165932574-oozie-admi-W]
# hadoop dfs -text /data/EDR/done.txt 2>/dev/null
```

2. Check the last timestamp for CubeExporter data cubes being exported.

```
[admin@collector-1 0000041-121020165932574-oozie-admi-W]
# hadoop dfs -text /data/CubeExporter/done.txt 2>/dev/null
```

3. Check the last timestamp for BulkStat Cubes being generated by the Bulk-Stat Job:

```
[admin@CISCO-COL1-147-11 ~]
# hadoop dfs -text /data/BulkStat/done.txt 2>/dev/null
# hadoop dfs -text /data/BSAgg15min/done.txt 2>/dev/null
```

4. Check the last timestamp for BulkStat data cubes being exported.

```
# hadoop dfs -text /data/BulkStatExporter_15min/done.txt
2>/dev/null
```

Data Validation on Insta Blades

1. SSH to the master Insta node and check the name of the database configured for EDR:

```
CISCOUCS-MUR-INSTA-01 [MUR-INSTA-CLUST: master] > en
```

```
CISCOUCS-MUR-INSTA-01 [MUR-INSTA-CLUST: master] # _shell
[admin@CISCOUCS-MUR-INSTA-01 ~]# cli -t "en" "conf t" "show runn
full" |
grep "insta instance 0 cubes-database" | awk -F ' ' '{print $5}'
DATABASE_MURAL
```

2. Connect to idbmysql and select the database:

```
[admin@CISCOUCS-MUR-INSTA-01 ~]# idbmysql
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 151
Server version: 5.1.39 MySQL Embedded / Calpont InfiniDB
Enterprise 2.2.9.1-1 GA (Commercial)
Type 'help;' or '\h' for help. Type '\c' to clear the current
input statement.
mysql>
mysql> use DATABASE_MURAL;
Database changed
```

3. Get the value for the **mints** and **maxts** field for -1 aggregation level and 60 minute bin class:

```
mysql> select * from bin_metatable;

+----------+---------------------+------------+------------+---------+
| binclass | aggregationinterval | mints      | maxts      | bintype |
+----------+---------------------+------------+------------+---------+
| 60min    |                  -1| 1350126000 | 1350594000 | NULL    |
| 60min    |               86400| 1350086400 | 1350432000 | NULL    |
| 60min    |              604800|          0 |          0 | NULL    |
| 60min    |             2419200|          0 |          0 | NULL    |
+----------+---------------------+------------+------------+---------+

4 rows in set (1.14 sec)

Press Ctrl+D to exit
mysql> Bye
```

4. Convert the date format: Run the date command with the value of **maxts** captured from the step above. The example shows the user has processed data from Oct 13 11:00 to Oct 18 21:00.

```
[admin@CISCOUCS-MUR-INSTA-01 ~]# date -d @1350126000 Sat Oct 13
```

 Copyright © 2014, Cisco Systems, Inc.

```
11:00:00 UTC 2012
[admin@CISCOUCS-MUR-INSTA-01 ~]# date -d @1350594000 Thu Oct 18
21:00:00 UTC 2012
```

Validate Bulk Stats Data on the Insta Blade

1. SSH to the master Insta node and check the name of the database configured for EDR:

```
CISCOUCS-MUR-INSTA-01 [MUR-INSTA-CLUST: master] > en
CISCOUCS-MUR-INSTA-01 [MUR-INSTA-CLUST: master] # _shell
[admin@CISCOUCS-MUR-INSTA-01 ~]# cli -t "en" "conf t" "show runn
full" |
grep "insta instance 1 cubes-database" | awk -F ' ' '{print $5}
'bulkstats
```

2. Connect to **idbmysql** and select the database:

```
[admin@CISCOUCS-MUR-INSTA-01 ~]# idbmysql
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 151
Server version: 5.1.39 MySQL Embedded / Calpont InfiniDB
Enterprise 2.2.9.1-1 GA (Commercial)
Type 'help;' or '\h' for help. Type '\c' to clear the current
input statement.
mysql>
mysql> use bulkStats;
Database changed
```

3. Select the data from the table. Run the following command to get the value for the **mints** and **maxts** field for 900 aggregation interval:

```
mysql> select * from bin_metatable;
```

```
+----------+---------------------+------------+------------+---------+
| binclass | aggregationinterval | mints      | maxts      | binType |
+----------+---------------------+------------+------------+---------+
| 5min     |                  -1 |          0 |          0 | NULL    |
| 5min     |                 900 | 1364713200 | 1367293500 | NULL    |
| 5min     |                3600 | 1364713200 | 1365004800 | NULL    |
| 5min     |               86400 | 1364688000 | 1364860800 | NULL    |
| 5min     |              604800 |          0 |          0 | NULL    |
| 5min     |             2419200 |          0 |          0 | NULL    |
+----------+---------------------+------------+------------+---------+

6 rows in set (12.18 sec)

mysql> quit
```

4. Convert the date format: Run the date command with the value of maxts (captured from the step above) for the row which shows **aggregationinterval** as 900. This shows we have processed data from Sun Mar 31 07:00:00 UTC 2013 to Tue Apr 30 03:45:00 UTC 2013.

```
[admin@CachingCompute-GMS-2 ~]# date -d@1367293500
Tue Apr 30 03:45:00 UTC 2013
[admin@CachingCompute-GMS-2 ~]# date -d@1364713200
Sun Mar 31 07:00:00 UTC 2013
```

Start UI Processes and Verify Data

These final steps include starting the UI processes and verifying the data on the UI. Ensure that the URL is set up in the DNS for the production system.

Start the Rubix Tomcat Instance on Both UI Nodes

Note: You should only start UI Tomcat instances after at least two hours of data has been pushed into the Insta node.

1. Log in to the master UI node and start the EDR process:

```
> en
# conf t
(config) # pm process rubix restart
(config) # rubix modify-app atlas enable
(config) # rubix modify-app atlas modify-instance 1 enable
```

 Copyright © 2014, Cisco Systems, Inc.

2. Log in to the standby UI node and start the EDR process:

```
> en
# conf t
(config) # pm process rubix restart
(config) # rubix modify-app atlas enable
(config) # rubix modify-app atlas modify-instance 1 enable
```

NOTE: Wait for two minutes before moving to the next step.

3. Login to Master UI node and start other processes:

```
(config) # rubix modify-app bulkstats enable
(config) # rubix modify-app bulkstats modify-instance 1 enable
```

NOTE: Wait for two minutes before moving to the next step.

```
(config) # rubix modify-app reportAtlas enable
(config) # rubix modify-app reportAtlas modify-instance 1 enable
```

NOTE: Wait for two minutes before moving to the next step.

```
(config) # rubix modify-app rge enable
(config) # rubix modify-app rge modify-instance 1 enable
```

NOTE: Wait for two minutes before moving to the next step.

4. Login to the Standby UI node and start the other processes:

```
(config) # rubix modify-app bulkstats enable
(config) # rubix modify-app bulkstats modify-instance 1 enable
```

NOTE: Wait for two minutes before moving to the next step.

```
(config) # rubix modify-app reportAtlas enable
(config) # rubix modify-app reportAtlas modify-instance 1 enable
```

NOTE: Wait for two minutes before moving to the next step.

```
(config) # rubix modify-app rge enable
(config) # rubix modify-app rge modify-instance 1 enable
```

5. Access the UIs by going to the URL https://*domainName*:8443/ through your browser.

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

For example:

https://demo.sanmateo.com:8443/
Username: **admin**
Password: **admin123**

Note: Since the common certificate Installation procedure is not finalized, you must click the same URL for the BulkStats and RGE ports once before opening the actual URL:

Click the following ports once and accept the certificates:

https://*domainName*:20443/
https://*domainName*:30443/

For example:

https://demo.sanmateo.com:20443/
https://demo.sanmateo.com:30443/

Note: Once the installation is completed, be sure to back up the configurations. Refer to the *Cisco MURAL Operations and Troubleshooting Guide* for more information.

 Copyright © 2014, Cisco Systems, Inc.

Mandatory Parameters for Incoming ASR Files

The following is the list of mandatory headers that need to be present in files coming from the ASR so that the MURAL system can deduce meaningful information.

Mandatory Attributes for Flow EDRs for MURAL

Flow EDR data sent by the ASR platform to the MURAL system must contain the following attributes:

- sn-start-time
- sn-end-time
- radius-calling-station-id
- sn-app-protocol
- p2p-protocol
- sn-server-port
- sn-volume-amt-ip-bytes-downlink
- sn-volume-amt-ip-pkts-uplink
- sn-volume-amt-ip-pkts-downlink
- sn-volume-amt-ip-bytes-uplink

Sample:

```
#sn-start-time,sn-end-time,radius-calling-station-id,sn-server-
port,sn-app-protocol,sn-volume-amt-ip-bytes-uplink,sn-volume-amt-
ip-bytes-downlink,sn-volume-amt-ip-pkts-uplink,sn-volume-amt-ip-
pkts-downlink,p2p-protocol,ip-server-ip-address,bearer-3gpp rat-
type,voip-duration,sn-direction,traffic-type,bearer-3gpp imei,
bearer-3gpp sgsn-address,bearer-ggsn-address,sn-flow-end-time,sn-
flow-start-time,radius-called-station-id,bearer-3gpp user-location-
information,sn-subscriber-port,ip-protocol,sn-rulebase,tcp-os-
signature,bearer-3gpp charging-id

1381518310,1381518337,1000000018,70000,29,20000,20000,182,36,iax,
27.9.126.155,2,1,FromMobile,,,2.2.2.1,27.23.157.2,1381518337,
1381518310,Sushfone-2,231-10-1073-10065,43769,1985,rb31,,2
```

Mandatory HTTP EDR Attributes for MURAL

HTTP EDR data sent to the MURAL system must contain the following attributes:

- sn-start-time
- sn-end-time
- transaction-downlink-packets
- transaction-uplink-packets
- transaction-downlink-bytes
- transaction-uplink-bytes
- http-content type
- radius-calling-station-id
- http-host

Sample:

```
#sn-start-time,sn-end-time,radius-calling-station-id,transaction-
uplink-bytes,transaction-downlink-bytes,ip-subscriber-ip-address,
ip-server-ip-address,http-host,http-content type,http-url,voip-
duration,traffic-type,transaction-downlink-packets,transaction-
uplink-packets,bearer-3gpp rat-type,radius-called-station-id,tcp-
os-signature,bearer-3gpp imei,http-request method,http-reply code,
http-user-agent

1381518310,1381518338,1000000019,15000,15000,1.1.1.1,27.2.248.155,
images.craigslist.org,image/png,images.craigslist.org,11,,60,1,1,
Sushfone-1,,,GET,506 Variant Also Negotiates,"Dalvik/1.6.0 (Linux;
U; Android 4.0.3; Galaxy Nexus Build/ICL53F)"
```

ASR-Side Configuration

The corresponding configuration on the side of the ASR platform is as follows:

```
    edr-format edr-flow-format
      attribute  sn-start-time format seconds priority 1
      rule-variable ip subscriber-ip-address  priority 4
      attribute sn-subscriber-port priority 6
      attribute  sn-start-time format seconds priority 10
```


Copyright © 2014, Cisco Systems, Inc.

```
    attribute sn-charge-volume ip pkts uplink priority 12
    attribute sn-charge-volume ip pkts downlink priority 13
    attribute sn-charge-volume ip bytes uplink priority 14
    attribute sn-charge-volume ip bytes downlink priority 15
    attribute sn-direction priority 21
    rule-variable bearer ggsn-address  priority 23
    rule-variable bearer 3gpp2 bsid  priority 24
    attribute  sn-flow-start-time format seconds priority 26
    attribute  sn-flow-end-time format seconds priority 27
    attribute sn-flow-id priority 28
    attribute sn-closure-reason priority 29
    attribute radius-calling-station-id priority 30
    rule-variable p2p protocol  priority 31
    rule-variable bearer 3gpp imsi  priority 35
    attribute radius-called-station-id priority 40
    rule-variable ip server-ip-address  priority 60
    attribute sn-server-port priority 70
    attribute sn-app-protocol priority 80
    attribute sn-parent-protocol priority 81
    rule-variable ip protocol  priority 82
    attribute sn-volume-amt ip bytes uplink priority 100
    attribute sn-volume-amt ip bytes downlink priority 110
    attribute sn-volume-amt ip pkts uplink priority 120
    attribute sn-volume-amt ip pkts downlink priority 130
    rule-variable bearer 3gpp charging-id  priority 140
    rule-variable bearer 3gpp imei  priority 141
    rule-variable bearer 3gpp rat-type  priority 142
    rule-variable bearer 3gpp user-location-information  priority 143
    rule-variable bearer 3gpp sgsn-address  priority 144
    rule-variable traffic-type priority 160
    rule-variable voip-duration priority 170
    attribute  sn-end-time format seconds priority 180
  #exit
  edr-format edr-http-format
    rule-variable ip subscriber-ip-address  priority 4
    attribute sn-charge-volume ip pkts uplink priority 8
    attribute sn-charge-volume ip pkts downlink priority 9
    attribute  sn-start-time format seconds priority 10
```

```
    attribute sn-charge-volume ip bytes downlink priority 11
    attribute sn-charge-volume ip bytes uplink priority 12
    attribute  sn-end-time format seconds priority 20
    attribute radius-calling-station-id priority 30
    attribute radius-called-station-id priority 40
    rule-variable ip server-ip-address  priority 50
    rule-variable http user-agent  priority 55
    rule-variable http host  priority 70
    rule-variable http content type  priority 80
    attribute transaction-downlink-bytes priority 90
    attribute transaction-uplink-bytes priority 100
    attribute transaction-downlink-packets priority 110
    attribute transaction-uplink-packets priority 120
    rule-variable bearer 3gpp charging-id  priority 160
  #exit
```


Copyright © 2014, Cisco Systems, Inc.

Glossary

A

ASR Cisco ASR 5000 or ASR 5500 Series platforms. The routing platform that pushes the data flows to the MURAL system.

B

Big data A collection of data that is too large to manage, store, or analyze by traditional methods.

C

Caching node Hosts the Rubix engine and data cache. The Rubix engine queries the Insta nodes constantly and when new data is available, it fetches it to store in the data cache, so that it can respond more quickly to requests from the UI engine. The Caching node is sometimes called the Rubix node. The Caching node uses N+1 redundancy in active-active mode.

CIQ An Excel spreadsheet containing a site survey that you completed before beginning the installation process.

Collector Cluster The cluster of nodes consisting of the Collector nodes in active/standby High Availability clustering.

Collector node Collects data streams pushed to the Reflex platform, interprets the exported flows, enriches them with static data, and assembles data sets. The Collector stores the raw data in the Hadoop file system (HDFS) and sends it to the Compute node. The Collector node cluster can have any number of servers, in pairs for master and standby and uses 1+1 redundancy (transparent failover between pairs of active-active nodes).

Compute cluster The cluster consisting of the master and standby Compute nodes.

Compute node	Analyzes and aggregates the data, creating *data cubes*. Data cube is a convenient shorthand for the data structure, which is actually multi-dimensional and not limited to the three dimensions of a cube. The Compute node cluster can have any number of servers, depending on the implementation, and uses N+1 redundancy.
Cube engine	The Cube engine (Rubix) and the RG (report generation) engine are hosted on the UI node. The Cube engine forwards requests from the UI engine to the Insta node. It also prefetches data and locally caches it so that if the requested data is in the local cache, it can return the response directly to the UI node without querying the Insta node. The RG engine serves as the HTTP request server.

D

Direct-Attached SAN	A feature on Cisco UCS that enables you to attach a fiber-channel SAN directly to the Fabric Interconnects.
Downlink Rate	The average bytes received by the mobile device from the Internet during a selected interval.
Downlink Tonnage	The total amount of data received by the mobile device from the Internet.
DPI	Deep packet inspection-an advanced form of packet filtering that makes it possible to examine more than just the packet header.

F

Fabric Interconnect	Part of the Cisco UCS platform that provides both network connectivity and management capabilities to all attached blades and chassis.
FC Switch Mode	Fibre channel switch mode, a mode on the Fabric Interconnects.
FCID	A 24-bit field used to route frames through a FC network.
flogi database	Fabric login database.

G

Copyright © 2014, Cisco Systems, Inc.

GMS node General Management System. This node provides centralized management of the Reflex platform nodes, such as remote manufacturing of blades, patch management, monitoring of all nodes and operations, and importing and running node configurations.

H

Hadoop Open-source software that supports running applications on large clusters of hardware. See http://hadoop.apache.org/

I

Image Comprised of the operating system plus the application bundle.

Insta node Stores and manages the processed data in a columnar database. Manages the Insta database, which stores processed data cubes, commonly three to six years' worth. The Insta node cluster has two servers with 1+1 redundancy.

M

Manufacture Manufacturing a blade is the installing the MURAL operating system on the blade.

MIBs A database used for managing the entities in a communications network; associated with the Simple Network Management Protocol (SNMP).

MURAL Cisco Mobility Unified Reporting and Analytics (MURAL) application, which provides Web-based reporting and analytics abilities for deep packet inspection (DPI) data emerging from the network.

mural.xml file Contains sample configuration settings for your setup, including a logical grouping of the nodes into two chassis and configuration settings for all of the system components. These settings are based on the details of your network that were supplied in the Customer Information Questionaire (CIQ).

N

NX-OS Cisco Nexus Operating System

O

Oozie A workflow scheduler designed by Apache to manage Hadoop jobs. Oozie is bundled on the system and hosted on the Collector nodes.

P

PGSQL disks On a MURAL installation, these are storage partitions for the Postgres database and will store Rubix-related data.

R

Rate (bps) Amount of data sent and received per second between the device and the Internet. (Bytes per Second – bps)

RG engine Report Generation engine, which serves as the HTTP request server. In Cisco MURAL, it is hosted on the same blade as the UI engine.

Rubix engine See Cube engine.

Rubix node See Caching node.

S

SAN Storage area network. A high-speed network of storage devices that also connects those storage devices with servers. It provides block-level storage that can be accessed by the applications running on any networked servers.

T

TAC Type Allocation Code – The initial 8 digit portion of the 15 digit IMEI code used to uniquely identify wireless devices

Tonnage (MB) Total volume amount of data sent and received (Megabytes –MB)

U

UCS UCS 5108 Blade Server, connected to UCS 6248 Fabric Interconnects, hosts the MURAL application.

UI node See Caching node.

Uplink Rate The average bytes sent from the mobile device out to the Internet during a selected interval.

 Copyright © 2014, Cisco Systems, Inc.

Uplink Tonnage	The total amount of data sent from the mobile device out to the Internet.

V

vHBA Initiator Group	vHBA initiator groups determine the fiber channel zoning configuration for all vHBAs in a service profile. Cisco UCS Manager does not include any default vHBA initiator groups. You must create vHBA initiator groups in any service profile that is to be assigned to servers included in a zone.
VSAN	Virtual storage area network (SAN). A collection of ports from a set of connected Fiber Channel switches that form a virtual fabric that interconnects hosts and storage devices. A VSAN allows you to have logically separate fabrics without having to have different physical fabrics.

W

WWN	World Wide Name. A unique identifier that is assigned to a manufacturer that is hard-coded into a Fiber Channel (FC) device. A unique 16-digit hexadecimal number, such as 21-00-00-30-D9-00-12-34. This identifier is hard-coded into every FC host bus adapter (HBA).
WWNN	World Wide Name Node. A world wide name assigned to a port in a Fibre Channel fabric, in this context, a port on a MURAL node.
WWPN	World Wide Name assigned to a port in a Fibre Channel fabric to uniquely identify it in the network.

Z

Zoneset	A container for one or more zones in the fabric. Zones need to be a member of a zoneset in order to be used. Only one zoneset can be active at one time in any given VSAN.