



Cisco Virtualized Infrastructure Manager Administrator Guide, Release 2.2.24

First Published: 2018-06-29

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 NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://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. (1721R)

© 2018 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Managing Cisco NFVI 1

Managing Cisco NFVI Pods	1
General Guidelines for Pod Management	2
Identifying the Install Directory	4
Managing Hosts in Cisco VIM or NFVI Pods	4
Recovering Cisco NFVI Pods	7
Managing Nova Compute Scheduler Filters and User Data	9
Monitoring Cisco NFVI Health with CloudPulse	10
Assessing Cisco NFVI Status with Cloud-Sanity	12
Service Catalog URL	14
Get Token from Keystone	14
Get Service Catalog URL for Cloudpulse	15
Cloudpulse APIs	15
List of Cloudpulse Tests	15
Get detailed result of 1 test	16
Get List of Tests Available	17
Schedule a manual cloudpulse test:	17
Remove the results of a test	18
Checking Network Connections	18
Enabling NFVBench Post Deployment	19
NFVBench Usage	22
NFVBench Command-Line Options	23
Control Plane Verification	23
Fixed Rate Run Test	24
Packet Sizes	24
NDR and PDR Test	24

- Multi-chain Test 24
- Multi-Flow Test 25
- External Chain Test 25
- NFVBench Result Generation and Storage 25
- Interpretation of Results 26
- Advanced Configuration 30
- Cisco VIM CLI 30
- NFVBench REST Interface 31
- Enabling or Disabling Autobackup of Management Node 31
- Forwarding ELK logs to External Syslog Server 32
- Adding and Reconfiguring VIM Administrators 32
- Reconfigure of Proxy Post Install 33
- Enabling Custom Policy for VNF Manager Post Install 34
- Updating Containers in a Running Cisco VIM Cloud 34
- Updating Cisco VIM Software Using a USB 35
- Updating Cisco VIM Software Using Network Installation 38
- Upgrading Containers in a Running Cisco VIM Cloud 39
- Upgrading VIM Software Using a USB 42
- Upgrading Cisco VIM Software Using Network Installation 45
- VM Resizing 45
- Nova Migrate 46

CHAPTER 2

Cisco VIM REST API 47

- Overview to Cisco VIM REST API 47
- Cisco VIM REST API Resources 48

CHAPTER 3

Monitoring Cisco NFVI Performance 83

- Logging and Monitoring in Cisco NFVI 83
- Displaying Cisco VIM Log Files Using the CLI 85
- Logging Into the Kibana Dashboard 88
- Rotation of the Cisco VIM Logs 97
- Network Performance Test with NFVBench 97

CHAPTER 4

Managing Cisco NFVI Security 99

Verifying Management Node Network Permissions	99
Verifying Management Node File Permissions	100
Viewing Administrator Access Attempts	100
Verifying SELinux	101
Validating Port Listening Services	101
Validating Non-Root Users for OpenStack Services	102
Verifying Password Strength	102
Reconfiguring Passwords and OpenStack Configurations	103
Enabling NFVIMON Post Pod Install	106
Reconfiguring CIMC Password on an Existing Install	108
Increasing Provider and Tenant VLAN Ranges	109
Fernet Key Operations	109
Managing Certificates	110
Reconfiguring TLS Certificates	110
Enabling Keystone v3 on an Existing Install	111
LDAP/AD support with Keystone v3	112

CHAPTER 5

Managing Cisco NFVI Storage	115
Cisco NFVI Storage Architecture	115
Verifying and Displaying Ceph Storage Pools	116
Checking the Storage Cluster Health	117
Checking Glance Connectivity	118
Verifying Glance and Ceph Monitor Keyrings	119
Verifying Glance Image ID on Ceph	120
Checking Cinder Connectivity	120
Verifying the Cinder and Ceph Monitor Keyrings	121
Verifying the Cinder Volume ID on Ceph	122
Checking Nova Connectivity	122
Verifying the Nova and Ceph Monitor Keyrings	123
Verifying Nova Instance ID	124
Displaying Docker Disk Space Usage	125
Reconfiguring SwiftStack Integration	125
Integrating SwiftStack over TLS	126
Cinder Volume Backup on SwiftStack	127

Reconfiguring Administrator Source Networks 127
 Password Reset for Cisco VIM Management Node 128

CHAPTER 6 **Overview to Cisco VIM Unified Management 129**
 Cisco VIM Unified Management Overview 129
 Cisco VIM Unified Management Admin UI Overview 131
 Cisco VIM Unified Management Pod UI Overview 131

CHAPTER 7 **Managing Cisco VIM through Unified Management 133**
 UI Administrators Privileges and Responsibilities 133
 Pod UI Privileges and Responsibilities 134
 Adding Cisco VIM Pod 134
 Deleting Pod from Cisco VIM Unified Management 135
 Context Switching Within Insight 136

CHAPTER 8 **Managing Blueprints 137**
 Blueprints 137
 Blueprint Activation 137
 Viewing Blueprint Details 138
 Creating a Blueprint Using Upload Functionality 138
 Activating a Blueprint in an Existing Pod with OpenStack Installed 139
 Blueprint Management 139
 Creating a Blueprint for B-Series Server Platform 142
 Creating a Blueprint for C-Series Server Platform 157
 Downloading Blueprint 175
 Validating Blueprint 175
 Managing Post Install Features 176
 Monitoring the Pod 176
 Cross Launching Horizon 177
 NFVI Monitoring 177
 Run VMTP 177
 Run CloudPulse 178
 Run Cloud Sanity 179
 Run NFV Bench 180

Reconfiguring CIMC Password Through Unified Management 180

CHAPTER 9**Managing Pod Through Cisco VIM Unified Management 183**

Managing Hardware 183

Searching Compute and Storage nodes 184

POD Management 184

Managing Storage Nodes 185

Adding Storage Node 186

Deleting Storage Node 187

Managing Compute Nodes 187

Adding Compute Node 188

Deleting Compute Node 190

Managing Control Nodes 190

Replacing Control Node 190

Power Management 191

Power On a Compute Node 191

Powering Off Compute Node 192

Searching Compute and Storage Nodes 193

Managing Software 194

Reconfigure Openstack Passwords 195

Reconfigure OpenStack Services, TLS Certificates, and ELK Configurations 195

Reconfiguring CIMC Password through Unified Management 196

Reconfigure Optional Services 196

Reconfiguring Optional Features Through Unified Management 198

Pod User Administration 208

Managing Roles 208

Managing Users 209

Revoke Users 209

Edit Users 210

Managing Root CA Certificate 210

CHAPTER 10**Day 2 Operations of Cisco VIM Insight 213**

Shutting Down Cisco VIM Unified Management 213

Restarting Cisco VIM Unified Management 214

Restoring VIM Insight	214
Reconfiguring VIM Unified Management	216
Reconfiguring Insight TLS Certificate	216
Reconfiguring 3rd Party TLS Certificate	216
Reconfiguring Self Signed TLS Certificate	218
Switch from Self Signed TLS Certificate to 3rd Party TLS Certificate	220
Reconfiguring Insight MySQL Database Password	221
System generated Insight DB password	222
User supplied Insight DB password	223
Reconfiguring Unified Management SMTP Server	224
Update VIM Insight	226
Update Scenarios	226
Update VIM Insight with Internet Access	226
VIM Insight without Internet Access	227
Rollback VIM Unified Management	229
Commit VIM Unified Management	230

CHAPTER 11	Overview to the Cisco Virtual Topology System	233
	Understanding Cisco VTS	233
	Cisco VTS Architecture Overview	234
	Virtual Topology Forwarder	235
	Overview to Cisco VTF and VPP	235
	VPP + VHOSTUSER	236
	Virtual Topology System High Availability	237

CHAPTER 12	Managing Backup and Restore Operations	239
	Managing Backup and Restore Operations	239
	Backing Up the Management Node	239
	Backup with Forwarding ELK Logs to External Syslog Server	241
	Backing Up VIM UM	241
	Autobackup Insight	241
	Back Up Insight at Default Back Up Location	243
	Backup Insight at user defined backup location	244
	Restoring the Management Node	245

Management Node Autobackup 247

CHAPTER 13**Troubleshooting 249**

Displaying Cisco NFVI Node Names and IP Addresses 249

Verifying Cisco NFVI Node Interface Configurations 250

Displaying Cisco NFVI Node Network Configuration Files 251

Viewing Cisco NFVI Node Interface Bond Configuration Files 252

Viewing Cisco NFVI Node Route Information 252

Viewing Linux Network Namespace Route Information 253

Prior to Remove Storage Operation 253

Troubleshooting Cisco NFVI 255

Managing CIMC and ISO Installation 255

Management Node Installation Fails 255

Configuring Boot Order 256

PXE Failure Issue During Baremetal Step 256

Connecting to Docker Container 259

Management Node Recovery Scenarios 260

Recovering Compute Node Scenario 269

Running the Cisco VIM Technical Support Tool 271

Tech-Support Configuration File 272

Tech-Support When Servers Are Offline 274

Disk-Maintenance Tool to Manage Physical Drives 275

OSD-Maintenance Tool 278

Utility to Resolve Cisco VIM Hardware Validation Failures 281

Command Usage 281

Examples of Command Usage 282

Cisco VIM Client Debug Option 283



CHAPTER 1

Managing Cisco NFVI

The following topics provide general management procedures that you can perform if your implementation is Cisco VIM by itself or is Cisco VIM and Cisco VIM Insight.

- [Managing Cisco NFVI Pods, on page 1](#)
- [Managing Nova Compute Scheduler Filters and User Data, on page 9](#)
- [Monitoring Cisco NFVI Health with CloudPulse, on page 10](#)
- [Assessing Cisco NFVI Status with Cloud-Sanity, on page 12](#)
- [Service Catalog URL, on page 14](#)
- [Checking Network Connections, on page 18](#)
- [Enabling NFVBench Post Deployment, on page 19](#)
- [NFVBench Usage, on page 22](#)
- [Enabling or Disabling Autobackup of Management Node, on page 31](#)
- [Forwarding ELK logs to External Syslog Server, on page 32](#)
- [Adding and Reconfiguring VIM Administrators, on page 32](#)
- [Reconfigure of Proxy Post Install, on page 33](#)
- [Enabling Custom Policy for VNF Manager Post Install, on page 34](#)
- [Updating Containers in a Running Cisco VIM Cloud, on page 34](#)
- [Updating Cisco VIM Software Using a USB, on page 35](#)
- [Updating Cisco VIM Software Using Network Installation, on page 38](#)
- [Upgrading Containers in a Running Cisco VIM Cloud, on page 39](#)
- [Upgrading VIM Software Using a USB, on page 42](#)
- [Upgrading Cisco VIM Software Using Network Installation, on page 45](#)
- [VM Resizing, on page 45](#)
- [Nova Migrate, on page 46](#)

Managing Cisco NFVI Pods

You can perform OpenStack management operations on Cisco NFVI pods including addition and removal of Cisco NFVI compute and Ceph nodes, and replacement of controller nodes. Each action is mutually exclusive. Only one pod management action can be performed at any time. Before you perform a pod action, verify that the following requirements are met:

- The node is part of an existing pod.

- The node information exists in the setup_data.yaml file, if the pod management task is removal or replacement of a node.
- The node information does not exist in the setup_data.yaml file, if the pod management task is to add a node.

To perform pod actions, see the [Managing Hosts in Cisco VIM or NFVI Pods](#), on page 4 section.

General Guidelines for Pod Management

The setup_data.yaml file is the only user-generated configuration file that is used to install and manage the cloud. While many instances of pod management indicates that the setup_data.yaml file is modified, the administrator does not update the system generated setup_data.yaml file directly.



Note To avoid translation errors, we recommend that you avoid copying and pasting commands from the documents to the Linux CLI.

To update the setup_data.yaml file, do the following:

1. Copy the setup data into a local directory:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
```

2. Update the setup data manually:

```
[root@mgmt1 ~]# vi my_setup_data.yaml (update the targeted fields for the setup_data)
```

3. Run the reconfiguration command:

```
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml>
<pod_management_action>
```

In Cisco VIM, you can edit and enable a selected set of options in the setup_data.yaml file using the reconfigure option. After installation, you can change the values of the feature parameters. Unless specified, Cisco VIM does not support unconfiguring of the feature.

The following table summarizes the list of features that can be reconfigured after the installation of the pod.

Features enabled after post-pod deployment	Comments
Optional OpenStack Services	<ul style="list-style-type: none"> • Heat: OpenStack Orchestration Program • Keystone v3: Pod running Keystone v2 can be migrated to Keystone v3 • LDAP: Works only with Keystone v3. Full or partial reconfiguration can be done. All attributes except domain is reconfigurable.

Features enabled after post-pod deployment	Comments
Pod Monitoring	<ul style="list-style-type: none"> NFVIMON: third party monitoring from host to service level; needs involvement and planning with Cisco Advance Services. It supports only Keystone v2 in VIM.
Export of EFK logs to External Syslog Server	Reduces single point of failure on management node and provides data aggregation.
NFS for Elasticsearch Snapshot	NFS mount point for Elastic-search snapshot is used so that the disk on management node does not get full.
Admin Source Networks	White list filter for accessing management node admin service.
NFVBench	Tool to help measure cloud performance. Management node needs a 10G Intel NIC (4x10G 710, or 2x10G 520 Intel NIC).
EFK settings	Enables you to set EFK rotation frequency and size.
OpenStack service password	Implemented for security reasons, so that OpenStack passwords can be reset on-demand.
CIMC Password Reconfigure Post Install	Implemented for security reasons, so that CIMC passwords for C-series pod, can be reset on-demand.
SwiftStack Post Install	Integration with third-party Object-Store. The SwiftStack Post Install feature works only with Keystone v2.
TENANT_VLAN_RANGES and PROVIDER_VLAN_RANGES	Ability to increase the tenant and provider VLAN ranges on a pod that is up and running. It gives customers flexibility in network planning.
Support of Multiple External Syslog Servers	Ability to offload the OpenStack logs to an external Syslog server post-install.
Replace of Failed APIC Hosts and add more leaf nodes	Ability to replace Failed APIC Hosts, and add more leaf nodes to increase the fabric influence.
Make Netapp block storage end point secure	Ability to move the Netapp block storage endpoint from Clear to TLS post-deployment
Auto-backup of Management Node	Ability to enable/disable auto-backup of Management Node. It is possible to unconfigure the Management Node.
VIM Admins	Ability to configure non-root VIM Administrators.
EXTERNAL_LB_VIP_FQDN	Ability to enable TLS on external_vip through FQDN.
EXTERNAL_LB_VIP_TLS	Ability to enable TLS on external_vip through an IP address.

Features enabled after post-pod deployment	Comments
http_proxy and/or https_proxy	Ability to reconfigure http and/or https proxy servers.
Admin Privileges for VNF Manager (ESC) from a tenant domain	Ability to enable admin privileges for VNF Manager (ESC) from a tenant domain.
SRIOV_CARD_TYPE	Mechanism to go back and forth between 2-X520 and 2-XL710 as an SRIOV option in Cisco VIC NIC settings through reconfiguration.
NETAPP	Migrate NETAPP transport protocol from http to https.

Identifying the Install Directory

If the administrator is using CLI to manage the pod, the administrator must know the directory where the pod is installed from (refer to installer directory). To identify the installer directory of a pod, execute the following commands:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# ls -lrt | grep openstack-configs
lrwxrwxrwx. 1 root root      38 Mar 12 21:33 openstack-configs ->
/root/installer-<tagid>/openstack-configs
```

From the output, you can understand that the OpenStack-configs is a symbolic link to the installer directory.

Verify that the REST API server is running from the same installer directory location, by executing the following command:

```
# cd installer-<tagid>/tools
#./restapi.py -a status
Status of the REST API Server:  active (running) since Thu 2016-08-18 09:15:39 UTC; 9h ago
REST API launch directory: /root/installer-<tagid>/
```

Managing Hosts in Cisco VIM or NFVI Pods

To perform actions on the pod, run the commands specified in the following table. If you log in as root, manually change the directory to /root/installer-xxx to get to the correct working directory for these Cisco NFVI pod commands.

Table 1: Cisco NFVI Pod Management

Action	Steps	Restrictions
Remove block_storage or compute node	<ol style="list-style-type: none"> Remove the node information from the ROLES and SERVERS section of the setup_data.yaml file for the specific node. Enter one of the following commands. For compute nodes: <pre>ciscovim remove-computes --setupfile ~/MyDir/my_setup_data.yaml <"compute-1,compute-2"></pre> For storage nodes: <pre>ciscovim remove-storage --setupfile ~/MyDir/my_setup_data.yaml <"storage-1"></pre> 	<p>You can remove multiple compute nodes and only one storage at a time;</p> <p>The pod must have a minimum of one compute and two storage nodes after the removal action.</p> <p>In Cisco VIM the number of ceph OSD nodes can vary from 3 to 20. You can remove one OSD node at a time as part of the pod management.</p> <p>Note On a micro-pod expanded with standalone computes, only the standalone compute nodes can be removed. Pod management operation for storage node is not supported for micro-pod. In hyper-converged mode, compute management operations are not supported for hyper-converged nodes.</p>
Add block_storage or compute node	<ol style="list-style-type: none"> Add the node information from the ROLES and SERVERS section of the setup_data.yaml file for the specific node. Enter one of the following commands. For compute nodes: <pre>ciscovim add-computes --setupfile ~/MyDir/my_setup_data.yaml <"compute-1,compute-2"></pre> For storage nodes: <pre>ciscovim add-storage --setupfile ~/MyDir/my_setup_data.yaml <"storage-1"></pre> 	<p>You can add multiple compute nodes and only one storage node at a time.</p> <p>The pod must have a minimum of one compute, and two storage nodes before the addition action.</p> <p>In Cisco VIM the number of ceph OSD nodes can vary from 3 to 20. You can add one OSD node at a time as part of the pod management.</p> <p>Note On a micro-pod expanded with standalone computes, only the standalone compute nodes can be added. Pod management operation for storage node is not supported for micro-pod. In hyper-converged mode, compute management operations are not supported for hyper-converged nodes.</p>

Action	Steps	Restrictions
Replace controller node	<ol style="list-style-type: none"> 1. If the controller node is in a UCS C-Series pod, update the CIMC info node in the SERVERS section of the setup_data.yaml file for the specific node 2. For B-series only update the blade and chassis info 3. Enter the following command: <pre> ciscovim replace-controller --setupfile ~/MyDir/my_setup_data.yaml <"control-1"> </pre> 	<p>You can replace only one controller node at a time. The pod can have a maximum of three controller nodes.</p> <p>In Cisco VIM the replace controller node operation is supported in micro-pod.</p> <p>Note While replacing the controller node, the IP address and hostname are reused. So, do not update any other controller information other than CIMC access for C-series, and blade and chassis information for B-series. For micro-pod, this operation is supported on the AIO (all in one) nodes.</p>

When you add a compute or storage node to a UCS C-Series pod, you can increase the management/provision address pool. Similarly, for a UCS B-Series pod, you can increase the Cisco IMC pool to provide routing space flexibility for pod networking. Along with server information, these are the only items you can change in the setup_data.yaml file after the pod is deployed. To make changes to the management or provisioning sections and/or CIMC (for UCS B-Series pods) network section, you must not change the existing address block as defined on day 0. You can add only to the existing information by adding new address pool block(s) of address pool as shown in the following example:

```

NETWORKING:
:
:

networks:
-
  vlan_id: 99
  subnet: 172.31.231.0/25
  gateway: 172.31.231.1
  ## 'pool' can be defined with single ip or a range of ip
  pool:
    - 172.31.231.2, 172.31.231.5 --> IP address pool on Day-0
    - 172.31.231.7 to 172.31.231.12 --> IP address pool ext. on Day-n
    - 172.31.231.20
  segments:
    ## CIMC IP allocation. Needs to be an external routable network
    - cimc
-
  vlan_id: 2001
  subnet: 192.168.11.0/25
  gateway: 192.168.11.1
  ## 'pool' can be defined with single ip or a range of ip
  pool:
    - 192.168.11.2 to 192.168.11.5 --> IP address pool on Day-0
    - 192.168.11.7 to 192.168.11.12 --> IP address pool on day-n
    - 192.168.11.20 --> IP address pool on day-n
  segments:
    ## management and provision goes together
    - management

```

```
- provision
:
:
```

The IP address pool is the only change allowed in the networking space of the specified networks management/provision and/or CIMC (for B-series). The overall network must have enough address space to accommodate for future enhancement on day-0. After making the changes to servers, roles, and the corresponding address pool, you can execute the add compute/storage CLI shown above to add new nodes to the pod.

Recovering Cisco NFVI Pods

This section describes the recovery processes for Cisco NFVI control node and the pod that is installed through Cisco VIM. For recovery to succeed, a full Cisco VIM installation must have occurred in the past, and recovery is caused by a failure of one or more of the controller services such as Rabbit MQ, MariaDB, and other services. The management node must be up and running and all the nodes must be accessible through SSH without passwords from the management node. You can also use this procedure to recover from a planned shutdown or accidental power outage.

Cisco VIM supports the following control node recovery command:

```
# ciscovim cluster-recovery
```

The control node recovers after the network partition is resolved.



Note It may be possible that database sync between controller nodes takes time, which can result in cluster-recovery failure. In that case, wait for some time for the database sync to complete and then re-run cluster-recovery.

To make sure Nova services are good across compute nodes, execute the following command:

```
# source /root/openstack-configs/openrc
# nova service-list
```

To check for the overall cloud status, execute the following:

```
# cd installer-<tagid>/tools
# ./cloud_sanity.py -c all
```

In case of a complete pod outage, you must follow a sequence of steps to bring the pod back. The first step is to bring up the management node, and check that the management node containers are up and running using the docker ps -a command. After you bring up the management node, bring up all the other pod nodes. Make sure every node is reachable through password-less SSH from the management node. Verify that no network IP changes have occurred. You can get the node SSH IP access information from /root/openstack-config/mercury_servers_info.

Execute the following command sequence:

- Check the setup_data.yaml file and runtime consistency on the management node:

```
# cd /root/installer-<tagid>/tools
# ciscovim run --perform 1,3 -y
```

- Execute the cloud sanity command:

```
# cd/root/installer-<tagid>/tools
# ./cloud_sanity.py -c all
```

- Check the status of the REST API server and the corresponding directory where it is running:

```
# cd/root/installer-<tagid>/tools
# ./restapi.py -a status
Status of the REST API Server: active (running) since Thu 2016-08-18 09:15:39 UTC; 9h
ago
REST API launch directory: /root/installer-<tagid>/
```

- If the REST API server is not running from the right installer directory, execute the following to get it running from the correct directory:

```
# cd/root/installer-<tagid>/tools
# ./restapi.py -a setup

Check if the REST API server is running from the correct target directory
# ./restapi.py -a status
Status of the REST API Server: active (running) since Thu 2016-08-18 09:15:39 UTC; 9h
ago
REST API launch directory: /root/new-installer-<tagid>/
```

- Verify Nova services are good across the compute nodes by executing the following command:

```
# source /root/openstack-configs/openrc
# nova service-list
```

If cloud-sanity fails, execute cluster-recovery (ciscovim cluster-recovery), then re-execute the cloud-sanity and nova service-list steps as listed above.

Recovery of compute and OSD nodes requires network connectivity and reboot so that they can be accessed using SSH without password from the management node.

To shutdown, bring the pod down in the following sequence:

1. Shut down all VMs, then all the compute nodes
2. Shut down all storage nodes serially
3. Shut down all controllers one at a time
4. Shut down the management node
5. Shut down the networking gears

Bring the nodes up in reverse order, that is, start with networking gears, then the management node, storage nodes, control nodes, and compute nodes. Make sure that each node type is completely booted up before you move on to the next node type.

Validate the Cisco API server by running the following command:

```
ciscovim run --perform 1,3 -y
```

Run the cluster recovery command to bring up the POD post power-outage

```
# help on sub-command
ciscovim help cluster-recovery

# execute cluster-recovery
ciscovim cluster-recovery
```

```
# execute docker cloudpulse check
# ensure all containers are up
cloudpulse run --name docker_check
```

Validate if all the VMs are up (not in shutdown state). If any of the VMs are in down state, bring them up using the Horizon dashboard.

Managing Nova Compute Scheduler Filters and User Data

OpenStack Nova is an OpenStack component that provides on-demand access to compute resources by provisioning large networks of virtual machines (VMs). In addition to the standard Nova filters, Cisco VIM supports the following additional scheduler filters:

- **ServerGroupAffinityFilter**—Ensures that an instance is scheduled onto a host from a set of group hosts. To use this filter, you must create a server group with an affinity policy and pass a scheduler hint using `group` as the key and the server group UUID as the value. Use the **nova** command-line tool and the **--hint** flag. For example:

```
$ nova server-group-create --policy affinity group-1
$ nova boot --image IMAGE_ID --flavor 1 --hint group=SERVER_GROUP_UUID server-1
```

- **ServerGroupAntiAffinityFilter**—Ensures that each group instance is on a different host. To use this filter, you must create a server group with an anti-affinity policy and pass a scheduler hint, using `group` as the key and the server group UUID as the value. Use the **nova** command-line tool and the **--hint** flag. For example:

```
$ nova server-group-create --policy anti-affinity group-1
$ nova boot --image IMAGE_ID --flavor 1 --hint group=SERVER_GROUP_UUID server-1
```

- **SameHostFilter**—Within an instance set, schedules one instance on the same host as another instance. To use this filter, pass a scheduler hint using **same_host** as the key and a list of instance UUIDs as the value. Use the **nova** command-line tool and the **--hint** flag. For example:

```
$ nova boot --image IMAGE_ID --flavor 1 --hint same_host=INSTANCE_ID server-1
```

- **DifferentHostFilter**—Within an instance set, schedules one instance on a different host than another instance. To use this filter, pass a scheduler hint using **different_host** as the key and a list of instance UUIDs as the value. The filter is the opposite of **SameHostFilter**. Use the **nova** command-line tool and the **--hint** flag. For example:

```
$ nova boot --image IMAGE_ID --flavor 1 --hint different_host=INSTANCE_ID server-1
```

In addition to scheduler filters, you can set up user data files for cloud application initializations. A user data file is a special key in the metadata service that holds a file that cloud-aware applications in the guest instance can access. For example, one application that uses user data is the cloud-init system, an open-source package that is available on various Linux distributions. The cloud-init system handles early cloud instance initializations. The typical use case is to pass a shell script or a configuration file as user data during the Nova boot, for example:

```
$ nova boot --image IMAGE_ID --flavor 1 --hint user-data FILE_LOC server-1
```

Monitoring Cisco NFVI Health with CloudPulse

You can query the state of various Cisco NFVI OpenStack endpoints using CloudPulse, an OpenStack health-checking tool. By default, the tool automatically polls OpenStack Cinder, Glance, Nova, Neutron, Keystone, Rabbit, Mariadb, and Ceph every four minutes. However, you can use a CLI REST API call from the management node to get the status of these services in real time. You can integrate the CloudPulse API into your applications and get the health of the OpenStack services on demand. You can find additional information about using CloudPulse in the following OpenStack sites:

- <https://wiki.openstack.org/wiki/Cloudpulse>
- <https://wiki.openstack.org/wiki/Cloudpulseclient>
- <https://wiki.openstack.org/wiki/Cloudpulse/DeveloperNotes>
- <https://wiki.openstack.org/wiki/Cloudpulse/OperatorTests>
- <https://wiki.openstack.org/wiki/Cloudpulse/APIDocs>

CloudPulse has two set of tests: `endpoint_scenario` (runs as a cron or manually) and `operator test` (run manually). The supported Cloudpulse tests groups include:

- `nova_endpoint`
- `neutron_endpoint`
- `keystone_endpoint`
- `glance_endpoint`
- `cinder_endpoint`

Operator tests include:

- `ceph_check`—Executes the command, "ceph -f json status" on the Ceph-mon nodes and parses the output. If the result of the output is not "HEALTH_OK" `ceph_check` the reports for an error.
- `docker_check`—Finds out if all the Docker containers are in the running state in all the nodes. It the report for an error if any containers are in the Exited state. It runs the command "docker ps -aq --filter 'status=exited'".
- `galera_check`—Executes the command, "mysql 'SHOW STATUS;'" on the controller nodes and displays the status.
- `node_check`—Checks if all the nodes in the system are up and online. It also compares the result of "nova hypervisor list" and finds out if all the computes are available.
- `rabbitmq_check`—Runs the command, "rabbitmqctl cluster_status" on the controller nodes and finds out if the rabbitmq cluster is in quorum. If nodes are offline in the cluster `rabbitmq_check` the report is considered as failed.

CloudPulse servers are installed in containers on all control nodes. The CloudPulse client is installed on the management node by the Cisco VIM installer. To execute CloudPulse, source the `openrc` file in the `openstack-configs` directory and execute the following:

```
[root@MercRegTBl openstack-configs]# cloudpulse --help
usage: cloudpulse [--version] [--debug] [--os-cache]
```

```

[--os-region-name <region-name>]
[--os-tenant-id <auth-tenant-id>]
[--service-type <service-type>]
[--endpoint-type <endpoint-type>]
[--cloudpulse-api-version <cloudpulse-api-ver>]
[--os-cacert <ca-certificate>] [--insecure]
[--bypass-url <bypass-url>] [--os-auth-system <auth-system>]
[--os-username <username>] [--os-password <password>]
[--os-tenant-name <tenant-name>] [--os-token <token>]
[--os-auth-url <auth-url>]
<subcommand> ...

```

To check results of periodic CloudPulse runs:

```

# cd /root/openstack-configs
# source openrc
# cloudpulse result
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| uuid | id | name | testtype | state |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| bf7fac70-7e46-4577-b339-b1535b6237e8 | 3788 | glance_endpoint | periodic | success |
| 1f575ad6-0679-4e5d-bc15-952bade09f19 | 3791 | nova_endpoint | periodic | success |

```

To view all CloudPulse tests:

```

# cd /root/openstack-configs
# source openrc
# cloudpulse result
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| uuid | id | name | testtype | state |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| bf7fac70-7e46-4577-b339-b1535b6237e8 | 3788 | glance_endpoint | periodic | success |
| 1f575ad6-0679-4e5d-bc15-952bade09f19 | 3791 | nova_endpoint | periodic | success |
| 765083d0-e000-4146-8235-ca106fa89864 | 3794 | neutron_endpoint | periodic | success |
| c1c8e3ea-29bf-4fa8-91dd-c13a31042114 | 3797 | cinder_endpoint | periodic | success |
| 04b0cb48-16a3-40d3-aa18-582b8d25e105 | 3800 | keystone_endpoint | periodic | success |
| db42185f-12d9-47ff-b2f9-4337744bf7e5 | 3803 | glance_endpoint | periodic | success |
| 90aa9e7c-99ea-4410-8516-1c08beb4144e | 3806 | nova_endpoint | periodic | success |
| d393a959-c727-4b5e-9893-e229efb88893 | 3809 | neutron_endpoint | periodic | success |
| 50c31b57-d4e6-4cf1-a461-8228fa7a9be1 | 3812 | cinder_endpoint | periodic | success |
| d1245146-2683-40da-b0e6-dbf56e5f4379 | 3815 | keystone_endpoint | periodic | success |
| ce8b9165-5f26-4610-963c-3ff12062a10a | 3818 | glance_endpoint | periodic | success |
| 6a727168-8d47-4a1d-8aa0-65b942898214 | 3821 | nova_endpoint | periodic | success |
| 6fbf48ad-d97f-4a41-be39-e04668a328fd | 3824 | neutron_endpoint | periodic | success |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

To run a CloudPulse test on demand:

```

# cd /root/openstack-configs
# source openrc
# cloudpulse run --name <test_name>
# cloudpulse run --all-tests
# cloudpulse run --all-endpoint-tests
# cloudpulse run --all-operator-tests

```

To run a specific CloudPulse test on demand:

```
# cloudpulse run --name neutron_endpoint
+-----+
| Property | Value |
+-----+
| name     | neutron_endpoint |
| created_at | 2016-03-29T02:20:16.840581+00:00 |
| updated_at | None |
| state    | scheduled |
| result   | NotYetRun |
| testtype | manual |
| id       | 3827 |
| uuid     | 5cc39fa8-826c-4a91-9514-6c6de050e503 |
+-----+
```

To show detailed results of a specific CloudPulse run:

```
#cloudpulse show 5cc39fa8-826c-4a91-9514-6c6de050e503
+-----+
| Property | Value |
+-----+
| name     | neutron_endpoint |
| created_at | 2016-03-29T02:20:16+00:00 |
| updated_at | 2016-03-29T02:20:41+00:00 |
| state    | success |
| result   | success |
| testtype | manual |
| id       | 3827 |
| uuid     | 5cc39fa8-826c-4a91-9514-6c6de050e503 |
+-----+
```

To see the CloudPulse options, source the openrc file in openstack-configs dir and execute:

```
#cloudpulse --help
```

The CloudPulse project has a RESTful Http service called the Openstack Health API. Through this API cloudpulse allows the user to list the cloudpulse tests, create new cloudpulse tests and see the results of the cloudpulse results.

The API calls described in this documentation require keystone authentication. We can use the keystone v2 or v3 version for the authentication. The corresponding configuration must be configured properly in the cloudpulse config in order that the cloudpulse can reach the v2 or the v3 keystone API.

The Identity service generates authentication tokens that permit access to the Cloudpulse REST APIs. Clients obtain this token and the URL endpoints for other service APIs by supplying their valid credentials to the authentication service. Each time you make a REST API request to Cloudpulse, you need to supply your authentication token in the X-Auth-Token request header.

Assessing Cisco NFVI Status with Cloud-Sanity

The cloud-sanity tool is designed to give you a quick overall status of the Pods health. Cloud-sanity can run tests on all node types in the Pod: management, control, compute, and ceph storage.

The following are test areas that are supported in cloud-sanity:

1. RAID Disk health checks.
2. Basic network connectivity between the management node and all other nodes in the Pod.

3. Mariadb cluster size.
4. RabbitMQ operation and status.
5. Nova service and hypervisor list.
6. CEPHMon operation and status.
7. CEPHOSD operation and status.

To run the cloud-sanity tool, log in to the management node and navigate to the tools folder for the installation currently running. Following are the cloud-sanity run options. Cloud-sanity can be run on all nodes or a particular target role.

Step 1 To run the cloud sanity, complete the following steps:

```
# cd /root/installer-<tag>/tools
# ./cloud_sanity.py -h
usage: cloud_sanity.py [-h] [--check CHECK] [--list] [--verbose]

cloud sanity helper

optional arguments:
  -h, --help            show this help message and exit
  --check CHECK, -c CHECK
                        all - Run all sanity checks. [default action]
                        control - Run controller sanity checks.
                        compute - Run compute sanity checks.
                        cephmon - Run cephmon sanity checks.
                        cephosd - Run cephosd sanity checks.
                        management - Run Management node sanity checks
  --list, -l            List all the available sanity checks.
  --verbose, -v        Run the sanity in verbose mode.
```

Step 2 To run all the cloud-sanity tests, select *all* as the check option.

```
./cloud_sanity.py --check all
Executing All Cloud Sanity in quiet mode. This takes some time.
Cloud Sanity Results
```

Role	Task	Result
Management	Management - Disk maintenance RAID Health *****	PASSED
Management	Management - Disk maintenance VD Health *****	PASSED
Control	Control - Ping All Controller Nodes *****	PASSED
Control	Control - Ping internal VIP *****	PASSED
Control	Control - Check Mariadb cluster size *****	PASSED
Control	Control - Check RabbitMQ is running *****	PASSED
Control	Control - Check RabbitMQ cluster status *****	PASSED
Control	Control - Check Nova service list *****	PASSED
Control	Control - Disk maintenance RAID Health *****	PASSED
Control	Control - Disk maintenance VD Health *****	PASSED

```

| Compute | Compute - Ping All Compute Nodes ***** | PASSED |
| Compute | Compute - Check Nova Hypervisor list ***** | PASSED |
| Compute | Compute - Disk maintenance RAID Health ***** | PASSED |
| Compute | Compute - Disk maintenance VD Health ***** | PASSED |
| CephMon | CephMon - Check cephmon is running ***** | PASSED |
| CephMon | CephMon - CEPH cluster check ***** | PASSED |
| CephMon | CephMon - Check Ceph Mon status ***** | PASSED |
| CephMon | CephMon - Check Ceph Mon results ***** | PASSED |
| CephOSD | CephOSD - Ping All Storage Nodes ***** | PASSED |
| CephOSD | CephOSD - Check OSD result with osdinfo ***** | PASSED |
| CephOSD | CephOSD - Check OSD result without osdinfo ***** | PASSED |
+-----+
[PASSED] Cloud Sanity All Checks Passed

```

The cloud-sanity tests use the disk-maintenance and osd-maintenance tools to assess overall health and status of RAID disks and OSD status.

Note Failures that are detected in RAID disk health and CEPHOSD operational status is evaluated with the disk-maintenance and osd-maintenance tools.

Service Catalog URL

The OpenStack Keystone service catalog allows API clients to dynamically discover and navigate to cloud services. Cloudpulse has its own service URL which is added to the Keystone service catalog. You need to send a token request to Keystone to find the service URL of cloudpulse. The token request lists all the catalog of services available.

Get Token from Keystone

To get the token from keystone run the following commands:

Resource URI

Verb	URI
POST	http://<controller_lb_ip>:5000/v2.0/tokens

Example

```

JSON Request
POST / v2.0/tokens
Accept: application/json
{

```

```

    "auth": {
      "passwordCredentials":{
        "username": "admin",
        "password": "iVP1YciVKoMGId1O"
      }
    }
  }
}

JSON Response
200 OK
Content-Type: application/json
{
  "access": {
    "token": {
      "issued_at": "2017-03-29T09:54:01.000000Z",
      "expires": "2017-03-29T10:54:01Z",
      "id":
"gAAAAABY24Q5TDIqizuGmhOXakV2rIzSvSPQpMAmC7SA2UzUXZQXSH-ME98d3Fp4FsJ16G561a420B4BK0fy1cykL22Eco9",
.....
.....
}
}

```

Get Service Catalog URL for Cloudpulse

Resource URI

Verb	URI
GET	http://<controller_ip>:35357/v2.0/endpoints

Example

```

JSON Request
GET /v2.0/endpoints
Accept: application/json

```

```

JSON Response
200 OK
Content-Type: application/json
{"endpoints": [
  {"internalurl": "http://<controller>:9999",
   "adminurl": "http://<controller>:9999",
   "publicurl": "http://<controller>:9999"
  }
]}
}

```

Cloudpulse APIs

The following are a list of APIs and the corresponding functions that the API performs. The cloudpulse APIs is accessed with the X-Auth-Token which contains the token which is received from the Keystone token generation API mentioned in the preceding panel.

List of Cloudpulse Tests

To get the list of cloudpulse tests:

Resource URI

Verb	URI
GET	http://<controller_ip>:9999/cpulse

Example

```
JSON Request
GET /cpulse
Accept: application/json
```

```
JSON Response
200 OK
Content-Type: application/json
{
  "cpulses": [
    {
      "name": "galera_check",
      "state": "success",
      "result": "ActiveNodes:16.0.0.37,16.0.0.17,16.0.0.27",
      "testtype": "periodic",
      "id": 4122,
      "uuid": "alb52d0a-ca72-448a-8cc0-5bf210438d89"
    }
  ]
}
```

Get detailed result of 1 test

To get detailed result of the test.

Resource URI

Verb	URI
GET	http://<controller_ip>:9999/cpulse/<uuid>

Uuid : uuid of the test

Example

```
JSON Request
GET /cpulse/e6d4de91-8311-4343-973b-c507d8806e94
Accept: application/json
```

```
JSON Response
200 OK
Content-Type: application/json
{
  "name": "galera_check",
  "state": "success",
  "result": "ActiveNodes:16.0.0.37,16.0.0.17,16.0.0.27",
  "testtype": "periodic",
  "id": 4122,
  "uuid": " e6d4de91-8311-4343-973b-c507d8806e94"
}
```

Get List of Tests Available

To get a list of available cloudpulse tests:

Resource URI

Verb	URI
GET	http://<controller_ip>:9999/cpulse/list_tests

Example

```
JSON Request
GET /cpulse/list_tests
Accept: application/json
```

```
JSON Response
200 OK
Content-Type: application/json
{
  "endpoint_scenario":
  "all_endpoint_tests\ncinder_endpoint\n glance_endpoint\nkeystone_endpoint\nneutron_endpoint\nnova_endpoint",

  "operator_scenario":
  "all_operator_tests\nceph_check\ndocker_check\n galera_check\nnode_check\nrabbitmq_check"
}
```

Schedule a manual cloudpulse test:

To schedule a manual test of cloudpulse run the following commands:

Resource URI

Verb	URI
POST	http://<controller_ip>:9999/cpulse

Example

```
JSON Request
POST /cpulse
Accept: application/json
{
  "name": "galera_check"
}
```

```
JSON Response
200 OK
Content-Type: application/json
{
  "name": "galera_check",
  "state": "scheduled",
  "result": "NotYetRun",
  "testtype": "manual",
  "id": 4122,
  "uuid": " e6d4de91-8311-4343-973b-c507d8806e94"
}
```

Remove the results of a test

To remove the results of a test.

Resource URI

Verb	URI
DELETE	http://<controller_ip>:9999/cpulse/<uuid>

Uuid : uuid of the test

Example

```
JSON Request
DELETE /cpulse/68ffaae3-9274-46fd-b52f-ba2d039c8654
Accept: application/json
```

```
JSON Response
204 No Content
```

Checking Network Connections

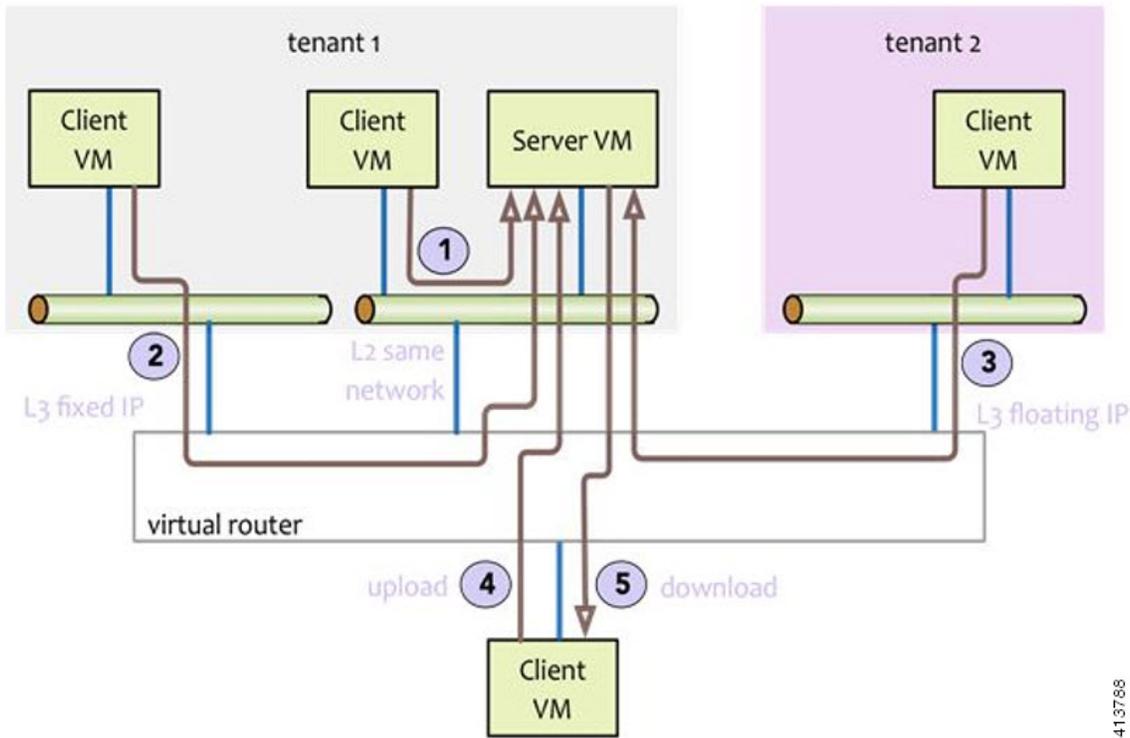
You can use Virtual Machine Through Put (VMTP) to check Layer 2 and Layer 3 data plane traffic between Cisco NFVI compute nodes. VMTP performs ping connectivity, round trip time measurement (latency), and TCP/UDP throughput measurement for the following Cisco NFVI east to west VM-to-VM flows:

- Same network (private fixed IP, flow number 1).
- Different network using fixed IP (same as intra-tenant L3 fixed IP, flow number 2).
- Different network using floating IP and NAT (same as floating IP inter-tenant L3, flow number 3.)
- When an external Linux host is available for testing north to south flows, external host to VM download and upload throughput and latency (L3/floating IP, flow numbers 4 and 5).

The following figure shows the traffic flows VMTP measures. Cloud traffic flows are checked during Cisco VIM installation and can be checked at any later time by entering the following command:

```
$ ciscovim run --perform 8 -y
```

Figure 1: VMTP Cloud Traffic Monitoring



413788

Enabling NFVBench Post Deployment

NFVBench is a data plane performance benchmark tool for NFVI that can be optionally installed after the pod deployment.

NFVBench is used to:

- Verify that the data plane is working properly and efficiently when using well defined packet paths that are typical of NFV service chains.
- Measure the actual performance of your data plane so that you can estimate what VNFs can expect from the infrastructure when it comes to receiving and sending packets.

While VMTP only measures VM to VM traffic, NFVBench measures traffic flowing from an integrated software traffic generator (TRex) running on the management node to the ToR switches to test VMs running in compute nodes.

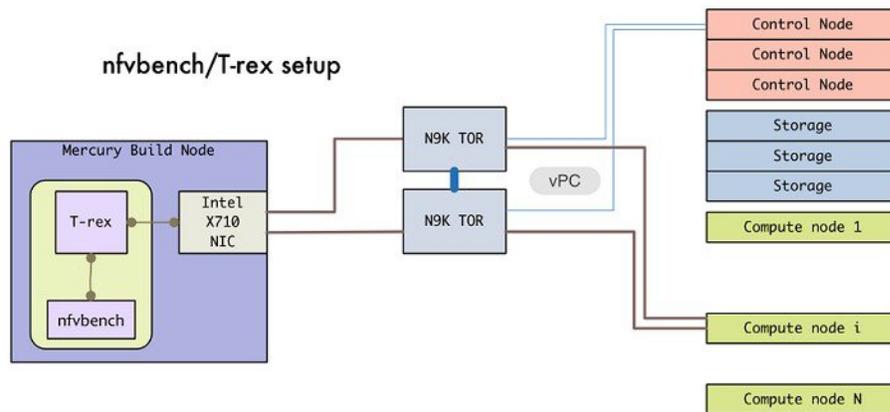
In Cisco VIM, the NFVBench (performance benchmark) is an optional tool. You can deploy NFVBench after the installation of the pod.

Before you begin

- A 10GE Intel NIC (Intel X710 NIC (4 x 10G)) must be installed on a management node.

- A TRex traffic generator which uses DPDK interface to interact with Intel NIC and makes use of hardware, instead of software to generate packets. This approach is more scalable and enables NFVBench to perform tests without software limitations.
- Wire two physical interfaces of the Intel NIC to the TOR switches (as shown in the following figure).

Figure 2: NFVBench topology setup



Procedure

	Command or Action	Purpose
Step 1	Enable the NFVBench configuration in the setup_data.yaml file.	<p>Sample configuration files for OVS/VLAN or VPP mechanism driver:</p> <pre>NFVBENCH: enabled: True # True or False tor_info: {TORa: eth1/42, TORb: eth1/42} # mandatory # tor_info: {TOR: 'eth1/42,eth1/43'} # use if there is only one TOR switch # nic_ports: 3,4 # Optional input, indicates which 2 of the 4 available ports # of 10G Intel NIC on the management node is NFVbench tool using # to send and receive traffic. # Defaults to the first 2 ports of NIC (ports 1 and 2) if not specified. # Port number must be between 1 and 4, one port cannot be used twice. # Example: # nic_ports: 1,4 # the first and the last port of Intel NIC are used</pre> <p>Sample configuration for VTS mechanism driver:</p> <pre>NFVBENCH: enabled: True # True or False tor_info: {TORa: eth1/42, TORb: eth1/42} # mandatory vtep_vlans: 1500,1501 # Mandatory and needed only for VTS/VXLAN. # Specify any pair of</pre>

	Command or Action	Purpose
		<pre> unused VLAN ids to be used # for VLAN to VxLAN encapsulation in TOR switch. # tor_info: {TOR: 'eth1/42,eth1/43'} # Use if there is only one TOR switch. # nic_ports: 3,4 # Optional input, indicates which 2 of the 4 available ports # of 10G Intel NIC on the management node is NFVbench tool using # to send and receive traffic. # Defaults to the first 2 ports of NIC (ports 1 and 2) if not specified. # Port number must be between 1 and 4, one port cannot be used twice. # Example: # nic_ports: 1,4 # the first and the last port of Intel NIC are used VTS_PARAMETERS: ... VTS_DAY0: '<True False>' # Required parameter when VTS enabled VTS_USERNAME: '<vts_username>' # Required parameter when VTS enabled VTS_PASSWORD: '<vts_password>' # Required parameter when VTS enabled VTS_NCS_IP: '11.11.11.111' # '<vts_ncs_ip>', mandatory when VTS enabled VTC_SSH_USERNAME: 'admin' # '<vtc_ssh_username>', mandatory for NFVbench VTC_SSH_PASSWORD: 'my_password' # '<vtc_ssh_password>', mandatory for NFVbench </pre>
<p>Step 2</p>	<p>Configuring minimal settings of NFVBench:</p>	<pre> # Minimal settings required for NFVbench TORSWITCHINFO: CONFIGURE_TORS: <True or False> # True if switches should be configured to support NFVbench ... SWITCHDETAILS: - hostname: 'TORa' # Hostname matching 'tor_info' switch name. username: 'admin' # Login username for switch user. password: 'my_password' # Login password for switch user. ssh_ip: '172.31.230.123' # SSH IP for switch. - hostname: 'TORb' username: 'admin' password: 'my_password' ssh_ip: '172.31.230.124' </pre> <p>TOR switches will be configured based on information provided in tor_info. Two ports specified by interfaces are configured in trunk mode. In order to access them and retrieve TX/RX counters you need the Login details for TOR switches. It is not required to set 'CONFIGURE_TORS' to 'True', but then manual configuration is necessary.</p>

	Command or Action	Purpose
		With VTS as mechanism driver additional settings are needed. NFVBench needs access to VTS NCS to perform cleanup after it detaches the traffic generator port from VTS. Also a pair of VTEP VLANs is required for VLAN to VxLAN mapping. Value can be any pair of unused VLAN ID.
Step 3	Reconfigure Cisco VIM to create a NFVBench container. To reconfigure add necessary configuration to the setup_data.yaml file, run the reconfigure command as follows.	<pre>[root@mgmt1 ~]# cd /root/ [root@mgmt1 ~]# mkdir MyDir [root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml /root/MyDir/ [root@mgmt1 ~]# cd /root/ [root@mgmt1 ~]# # update the setup_data to include NFVBENCH section [root@mgmt1 ~]# cd /root/MyDir/ [root@mgmt1 ~]# vi setup_data.yaml [root@mgmt1 ~]# cd ~/installer-xxxx [root@mgmt1 ~]# ciscovim --setupfile /root/MyDir/setup_data.yaml reconfigure</pre> <p>After reconfiguration is done, you can see NFVBench container up and is ready to use.</p>

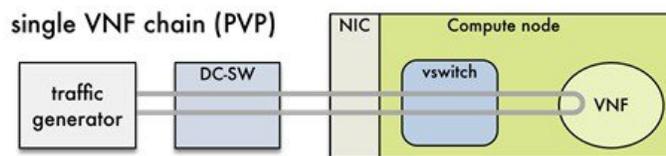
NFVBench Usage

Built-in packet paths

NFVBench can setup and stage three different packet paths.

The default packet path is called **PVP** (Physical - VM - Physical) and represents a typical service chain made of 1 VNF/VM:

Figure 3: Single VNF chain (PVP)



The traffic generator runs inside the NFVBench container on the management node. DC-SW represents the top of rack switch(es). The VNF is a test VM that contains a fast L3 router based on FD.io VPP. This VNF image can also be configured to run an L2 forwarder based on DPDK testpmd (both options generally yield roughly similar throughput results).

Traffic is made of UDP packets generated on the 2 physical interfaces (making it a bi-directional traffic). Packets are forwarded by the switch to the appropriate compute node before arriving to the virtual switch, then to the VNF before looping back to the traffic generator on the other interface. Proper stitching of the traffic on the switch is performed by NFVBench by using the appropriate mechanism (VLAN tagging for VLAN based deployments, VxLAN VTEP in the case of VTS deployments).

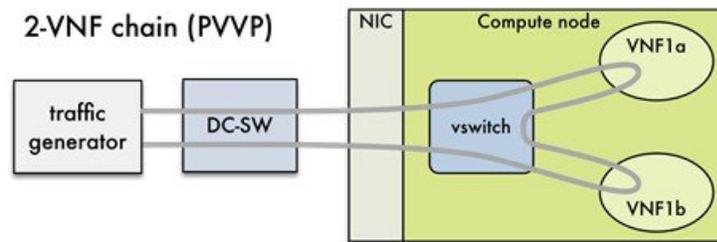
The performance of the PVP packet path provides a very good indication of the capabilities and efficiency of the NFVi data plane in the case of a single service chain made of 1 VNF/VM.

NFVBench also supports more complex service chains made of 2 VM in sequence and called PVVP (Physical-VM-VM-Physical).

In a PVVP packet path, the 2 VMs can reside on the same compute node (PVVP intra-node) or on different compute nodes (PVVP inter-node).

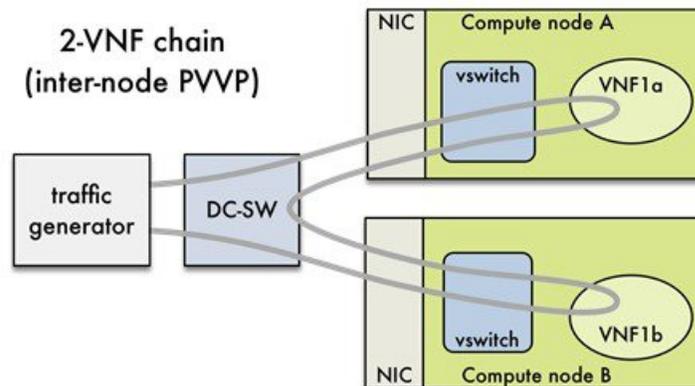
PVVP intra-node is more efficient when a virtual switch is used as packets do not have to go through the switch between the 2 VMs:

Figure 4: 2-VNF chain (PVVP)



PVVP inter-node requires packets to go through the switch and back between the 2 VMs.

Figure 5: 2-VNF chain(inter-node PVVP)



NFVBench Command-Line Options

The common NFVBench command-line options are displayed using the --help option:

```
[root@mgmt1 ~]# nfvbench --help
```

Control Plane Verification

If you are trying NFVBench for the first time, verify that the tool can stage the default packet path properly without sending any traffic.

The `--no-traffic` option exercises the control plane by creating a single test service chain with one VM, but does not send any traffic.

The following command stages only the default PVP packet path (but does not generate any traffic):

```
[root@mgmt1 ~]# nfvbench --no-traffic
```

Fixed Rate Run Test

The data plane traffic test is to generate traffic at a fixed rate for a fixed duration. For example, to generate a total of 10,000 packets per second (which is 5,000 packets per second per direction) for the default duration (60 seconds) and using the default frame size of 64 bytes:

```
[root@mgmt1 ~]# nfvbench --help
```

Packet Sizes

You can specify any list of frame sizes using the `-frame-size` option (pass as many as desired), including IMIX.

Following is an example, to run a fixed rate with IMIX and 1518 byte frames:

```
[root@mgmt1 ~]# nfvbench --rate 10kpps -frame-size IMIX -frame-size 1518
```

NDR and PDR Test

NDR and PDR test is used to determine the performance of the data plane in terms of throughput at a given drop rate.

- No Drop Rate(NDR)- It is the highest throughput achieved while allowing zero packet drop (allows a very low drop rate usually lesser than 0.001%).
- Partial Drop Rate (PDR)-It is the highest throughput achieved while allowing most at a given drop rate (typically less than 0.1%).

NDR is always less or equal to PDR.

To calculate the NDR and PDR for your pod run the following command:

```
[root@mgmt1 ~]# nfvbench --rate ndr_pdr
```

Multi-chain Test

In multi-chain test, each chain represents an independent packet path symbolizing real VNF chain. You can run multiple concurrent chains and better simulate network conditions in real production environment. Results with single chain versus with multiple chains usually vary because of services competing for resources (RAM, CPU, and network).

To stage and measure multiple service chains at the same time, use `--service-chain-count` flag or shorter `-scc` version.

The following example shows how to run the fixed rate run test with ten PVP chains:

```
[root@mgmt1 ~]# nfvbench -scc 10 --rate 100kpps
```

The following example shows how to run the NDR/PDR test with ten PVP chains:

```
[root@mgmt1 ~]# nfvbench -scc 10 --rate ndr_pdr
```

Multi-Flow Test

In Multi-flow test, one flow is defined by a source and destination MAC/IP/port tuple in the generated packets. It is possible to have many flows per chain. The maximum number of flows that are supported is in the order of 1 million flows per direction.

The following command runs three chains with a total of 100K flows per direction (for all chains):

```
[root@mgmt1 ~]# nfvsbench -scc 3 -fc 100k
```

External Chain Test

NFVBench measures the performance of chains that are pre-staged (using any means external to NFVBench). Such chains can be real VNFs with L3 routing capabilities or L2 forwarding chains.

This test is used when you want to use NFVBench for only traffic generation. In this case, NFVBench sends traffic from traffic generator and reports results without performing any configuration.

Do necessary configurations such as creating networks and VMs with a configuration that allows generated traffic to pass. NFVBench has to know the 2 edge networks to which the traffic generators are attached.

If the external chains only support L2 forwarding, the NFVBench configuration must:

- Enable VLAN tagging and define the VLAN IDs to use - if applicable (or disable vlan tagging if it is not required).
- The destination MAC to use in each direction (depends on the L2 forwarding mode in place in the service chain).

If the external chains support IPv4 routing, the NFVBench configuration must:

- Define the public IP addresses of the service chain end points (gateway IP) that used to discover destination MAC using ARP.
- Set the vlan tagging appropriately.

To measure performance for external chains, use the `--service-chain EXT` (or `-sc EXT`) option:

```
[root@mgmt1 ~]# nfvsbench -sc EXT
```



Note NFVBench cannot access ToR switches or v-switch in compute node.

NFVBench Result Generation and Storage

NFVBench detailed results can be stored in JSON format if you pass the `--json` option with a destination file name or the `--std-json` option with a destination folder pathname (if you want to use a standard file name generated by NFVBench). It is also possible to use both methods to generate the output into two different files at the same time:

```
[root@mgmt1 ~]# nfvsbench -scc 3 -fc 10 -fs 64 --json /tmp/nfvsbench/my.json --std-json /tmp/nfvsbench
```

The above command creates two JSON files in `/tmp/nfvsbench` container directory, which is mapped to the host directory. The first file is named `my.json`.

With the `--std-json` option, the standard NFVBench filename format follows this pattern:

```
<service-chain-type>-<service-chain-count>-<flow-count>-<frame-sizes>.json
```

Default chain is PVP and flag `-fs` was used to override traffic profile in the configuration file. With three chains and 10 flows specified file, the `PVP-3-10-64.json` is created.

Interpretation of Results

NFVBench prints data to the command line prompt in a table form. The data includes the current configuration, test devices details, and results computed based on traffic statistics.

Fixed Rate

Run the following command on NFVBench to view the traffic generated at fixed rate at different components of the packet path:

```
[root@mgmt1 ~]# nfvbench --rate 5kpps -fs IMIX
```

NFVBench summary consists of multiple blocks. In some cases, NFVBench displays lesser data. For example, in the output, the EXT chain does not access some path components (like switch). Therefore, the summary does not display.

```
===== NFVBench Summary ===== Date: 2017-03-28 19:59:53
NFVBench version 0.3.5 Openstack Neutron:
vSwitch: VTS Encapsulation: VxLAN
Benchmarks:
> Networks:
> Components:
> TOR:
Type: N9K Version:
10.28.108.249:
BIOS: 07.34
NXOS: 7.0(3)I2(2b) 10.28.108.248:
BIOS: 07.34
NXOS: 7.0(3)I2(2b)
> VTC:
Version:
build_date: 2017-03-03-05-41
build_number: 14
git_revision: 0983910
vts_version: 2.3.0.40 git_branch: vts231newton
job_name: vts231newton_gerrit_nightly
> Traffic Generator: Profile: trex-local Tool: TRex# Version:
build_date: Feb 16 2017 version: v2.18 built_by: hhaim build_time: 18:59:02
> Service chain:
> PVP:
> Traffic:
VPP version:
sjc04-pod3-compute-4: v17.04-rc0-98-g8bf68e8 Profile: custom_traffic_profile Bidirectional:
True
Flow count: 1
Service chains count: 1
Compute nodes: [u'nova:sjc04-pod3-compute-4'] Run Summary:

+-----+-----+-----+-----+-----+
| L2 Frame Size | Drop Rate | Avg Latency (usec) | Min Latency (usec) |
| Max Latency (usec) |
+-----+-----+-----+-----+-----+
| IMIX | 0.0000% | 16.50 | 10.00 |
| 241.00 |
+-----+-----+-----+-----+-----+
```

L2 frame size: IMIX
 Chain analysis duration: 73 seconds

Run Config:

```

+-----+-----+-----+-----+
| Direction | Duration (sec) | Rate | Rate |
+-----+-----+-----+-----+
| Forward | 60 | 7.6367 Mbps | 2,500 pps |
+-----+-----+-----+-----+
| Reverse | 60 | 7.6367 Mbps | 2,500 pps |
+-----+-----+-----+-----+
| Total | 60 | 15.2733 Mbps | 5,000 pps |
+-----+-----+-----+-----+
    
```

Chain Analysis:

Interface	Device	Packets (fwd)	Drops (fwd)	Drop% (fwd)	Packets (rev)	Drops (rev)	Drop% (rev)
traffic-generator	trex	150,042	0	0.0000%			
vni-5098	n9k	150,042	0	0.0000%	150,042	0	0.0000%
vxlan_tunnel0	vpp	150,042	0	0.0000%	150,042	0	0.0000%
VirtualEthernet0/0/1	vpp	150,042	0	0.0000%	150,042	0	0.0000%
VirtualEthernet0/0/0	vpp	150,042	0	0.0000%	150,042	0	0.0000%
vxlan_tunnel1	vpp	150,042	0	0.0000%	150,042	0	0.0000%
vni-5099	n9k	150,042	0	0.0000%	150,042	0	0.0000%
traffic-generator	trex	150,042	0	0.0000%	150,042	0	0.0000%

```
Run as:
nfvbench -c /tmp/nfvbench/nfvbench.cfg --rate 5kpps -fs IMIX
```

Summary Interpretation:

Lines 1-34: General information about host system and used components.

Lines 35-45: Test-specific information about service chain, traffic profile, and compute nodes.

Lines 46-53: Summary of traffic profile run with results. A new row is added to the table for every packet size in test. The output displays the run summary for the IMIX packet size, but lines for 64B and 1518B can also be present. The Table contains following columns:

The run summary table includes the following columns:

- Drop Rate: The percentage of total drop rate for all chains and flows from the total traffic sent.
- Avg Latency: Average latency of average chain latencies in microseconds.
- Min Latency: Minimum latency of all chain latencies in microseconds.
- Max Latency: Maximum latency of all chain latencies in microseconds.

Lines 54-68: Length of the

Lines 69-89: Detailed analysis of test. Each row represents one interface on packet path in order they are visited. Left side of the table is for forward direction (from traffic generator port 0 to port 1), and the right side is for reverse direction (from traffic generator port 1 to port 0).

The chain analysis table has following columns:

- Interface: Interface name on devices in packet path.
- Device: Device name on which the interface is displayed in the first column is available.
- Packets (fwd): RX counter on given interface, only the first row is TX counter (it is the beginning of packet path).
- Drops (fwd): Amount of packets being dropped between current and previous interface.
- Drop% (fwd): Percentage of dropped packets on this interface to the total packet drops.
- Packets (rev): Similar to Packets (fwd) but for the reverse direction.
- Drops (rev): Similar to Drops (fwd) but for the reverse direction.
- Drop% (rev): Similar to Drop% (fwd) but for reverse direction.

This type of summary is very useful for finding bottlenecks or to verify if the system can handle certain fixed rate of traffic.

NDR/PDR

The test result shows throughput values in different units for both NDR and PDR with latency statistics (minimum, maximum, average) for each test.

```
[root@mgmt1 ~]# nfvbench --rate ndr_pdr -fs IMIX
===== NFVBench Summary =====
Date: 2017-03-28 20:20:46
NFVBench version 0.3.5 Openstack Neutron:
vSwitch: VTS Encapsulation: VxLAN
Benchmarks:
> Networks:
> Components:
> TOR:
Type: N9K Version:
10.28.108.249:
BIOS: 07.34
NXOS: 7.0(3)I2(2b) 10.28.108.248:
```

BIOS: 07.34

> VTC:

NXOS: 7.0(3)I2(2b)

```
Version:
build_date: 2017-03-03-05-41
build_number: 14
git_revision: 0983910
vts_version: 2.3.0.40 git_branch: vts231newton
job_name: vts231newton_gerrit_nightly
> Traffic Generator: Profile: trex-local Tool: TRex Version:
build_date: Feb 16 2017 version: v2.18 built_by: hhaim build_time: 18:59:02
> Measurement Parameters: NDR: 0.001
PDR: 0.1
> Service chain:
> PVP:
> Traffic:
VPP version:
sjc04-pod3-compute-4: v17.04-rc0~98-g8bf68e8 Profile: custom_traffic_profile Bidirectional:
True
Flow count: 1
Service chains count: 1
Compute nodes: [u'nova:sjc04-pod3-compute-4'] Run Summary:
```

	L2 Frame Size	Rate (fwd+rev)	Rate (fwd+rev)	Avg Drop Rate	Avg Latency (usec)	Min Latency (usec)	Max Latency (usec)
NDR	IMIX	4.5703 Gbps	1,496,173 pps	0.0006%	131.33	10.00	404.00
PDR	IMIX	4.7168 Gbps	1,544,128 pps	0.0553%	205.72	20.00	733.00

L2 frame size: IMIX Chain analysis duration: 961 seconds NDR search duration: 661 seconds
PDR search duration: 300 seconds

Run Config:

Direction	Duration (sec)	Rate	Rate
Forward	60	2.3584 Gbps	772,064 pps
Reverse	60	2.3584 Gbps	772,064 pps
Total	60	4.7168 Gbps	1,544,128 pps

Lines 1-48: Similar to the fixed rate run output explained above.


```

| Name          | Value          |
+-----+-----+
| status        | not_run       |
| nfvsbench_request | {"rate": "5kpps"} |
| uuid         | 0f131259-d20f-420f-840d-363bdcc26eb9 |
| created_at   | 2017-06-26T18:15:24.228637 |
+-----+-----+

```

Run the following command with the returned UUID to poll status:

```

[root@mgmt1 ~]# ciscovim nfvsbench --stat 0f131259-d20f-420f-840d-363bdcc26eb9
+-----+-----+
| Name          | Value          |
+-----+-----+
| status        | nfvsbench_running |
| nfvsbench_request | {"rate": "5kpps"} |
| uuid         | 0f131259-d20f-420f-840d-363bdcc26eb9 |
| created_at   | 2017-06-26T18:15:24.228637 |
| updated_at   | 2017-06-26T18:15:32.385080 |
+-----+-----+

```

```

+-----+-----+
| Name          | Value          |
+-----+-----+
| status        | nfvsbench_completed |
| nfvsbench_request | {"rate": "5kpps"} |
| uuid         | 0f131259-d20f-420f-840d-363bdcc26eb9 |
| created_at   | 2017-06-26T18:15:24.228637 |
| updated_at   | 2017-06-26T18:18:32.045616 |
+-----+-----+

```

When the test is done, retrieve results in a JSON format:

```

[root@mgmt1 ~]# ciscovim nfvsbench --json 0f131259-d20f-420f-840d-363bdcc26eb9
{"status": "PROCESSED", "message": {"date": "2017-06-26 11:15:37", ...}}

```

NFVBench REST Interface

When enabled, the NFVBench container can also take benchmark request from a local REST interface. Access is only local to the management node in the current Cisco VIM version (that is the REST client must run on the management node).

Details on the REST interface calls can be found in Chapter 2, Cisco VIM REST API Resources.

Enabling or Disabling Autobackup of Management Node

Cisco VIM supports the backup and recovery of the management node. By default, the feature is enabled. Auto snapshot of the management node happens during pod management operation. You can disable the auto backup of the management node.

To enable or disable the management node, update the `setup_data.yaml` file as follows:

```

# AutoBackup Configuration
# Default is True
#autobackup: <True or False>

```

Take a backup of `setupdata` file and update it manually with the configuration details by running the following command:

```

[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir

```

```
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml /root/MyDir/
[root@mgmt1 ~]# # update the setup_data to change autobackup
[root@mgmt1 ~]# cd /root/MyDir/
[root@mgmt1 ~]# vi setup_data.yaml
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile /root/MyDir/setup_data.yaml reconfigure
```

Forwarding ELK logs to External Syslog Server

Cisco VIM supports backup and recovery of the management node. To keep the process predictable and to avoid loss of logs, the software supports the capability of forwarding the ELK logs to multiple external syslog servers (Minimum 1 and Maximum 3). The capability is introduced to enable this feature after the pod is up and running, with Cisco VIM, through the reconfigure option.

The Syslog Export reconfigure option supports the following options:

- Enable forwarding of ELK logs to External Syslog Server on a pod that is already up and running.
- Reconfigure existing External Syslog Setting to point to a different syslog cluster.

The following section needs to be configured in the setup_data.yaml file.

```
#####
## SYSLOG EXPORT SETTINGS
#####
SYSLOG_EXPORT_SETTINGS:
-
  remote_host: <Syslog_ipv4_or_v6_addr> # requiredIP address of the remote syslog
  server protocol : udp # defaults to udp
  facility : <string> # required; possible values local[0-7]or user
  severity : <string; suggested value: debug>
  port : <int>; # defaults, port number to 514
  clients : 'ELK' # defaults and restricted to ELK;
```

Take a backup of the setupdata file and update the file manually with the configs listed in the preceding section, then run the reconfigure command as follows:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml /root/MyDir/
[root@mgmt1 ~]# # update the setup_data to include Syslog Export info
[root@mgmt1 ~]# cd /root/MyDir/
[root@mgmt1 ~]# vi setup_data.yaml
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile /root/MyDir/setup_data.yaml reconfigure
```

With this configuration, you should now be able to use export ELK logs to an external syslog server. On the remote host, verify if the logs are forwarded from the management node.

Adding and Reconfiguring VIM Administrators

Cisco VIM supports management of the VIM Administrators. VIM administrator has the permission to log in to the management node through SSH or the console using the configured password. By configuring to one

VIM admin account, administrators do not have to share credentials. Administrators have individual accountability.

To enable one or more VIM administrators, perform the following steps:

Step 1 Take a backup of the setupdata file and update the file manually with the configurations listed as,

```
vim_admins:
- vim_admin_username: <username>
  vim_admin_password_hash: <sha512-password-hash>
- vim_admin_username: <username>
  vim_admin_password_hash: <sha512-password-hash>
- vim_admin_username: <username>
  vim_admin_password_hash: <sha512-password-hash>
```

The value of password hash must be in the standard sha512 format. # To generate the hash admin_password_hash should be the output from on the management node
python -c "import crypt; print crypt.crypt('<plaintext password>')"

Step 2 Run the reconfigure commands as follows:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml /root/MyDir/

# update the setup_data to include vim_admin info
[root@mgmt1 ~]# cd /root/MyDir/
[root@mgmt1 ~]# vi setup_data.yaml
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile /root/MyDir/setup_data.yaml reconfigure
```

Reconfigure of Proxy Post Install

During post-install you can update the http/https proxy server information that is listed in NETWORKING section of the setup_data.yaml.

To update the proxy in the post-VIM install follow these steps:

Step 1 Take a backup of the setupdata file and update the file manually with the configs listed as,

```
http_proxy_server: <a.b.c.d:port> # optional, needed if install is through internet, and the pod is
  behind a proxy
and/or
https_proxy_server: <a.b.c.d:port> # optional, needed if install is through internet, and the pod is
  behind a proxy
```

Step 2 Run the following command to reconfigure:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml /root/MyDir/

# update the setup_data to update the proxy info
[root@mgmt1 ~]# cd /root/MyDir/
[root@mgmt1 ~]# vi setup_data.yaml
```

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile /root/MyDir/setup_data.yaml reconfigure
```

Enabling Custom Policy for VNF Manager Post Install

During the post-installation of a cloud, Cisco VIM helps to enable a VNF Manager (such as ESC) to operate and manage tenant VMs in the OpenStack cloud, with additional privileged features.

Following are the steps to enable the custom policy for VNF Manager:

Step 1 Take a backup of the setupdata file and update the file manually with the configurations listed as,

```
ENABLE_ESC_PROV: True
```

Step 2 Run the following commands to reconfigure:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml /root/MyDir/

# update the setup_data to update the proxy info
[root@mgmt1 ~]# cd /root/MyDir/
[root@mgmt1 ~]# vi setup_data.yaml
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile /root/MyDir/setup_data.yaml reconfigure
```

Updating Containers in a Running Cisco VIM Cloud

Cisco VIM allows you to update all OpenStack and infrastructure services such as RabbitMQ, MariaDB, HAProxy, and management node containers such as Cobbler, ELK, VMTP and repo containers with almost no impact to the Cisco NFVI implementation. Updates allows you to integrate Cisco VIM patch releases without redeploying the Cisco NFVI stack from the beginning. Updates have minimal service impact because they run serially component by component one node at a time. If an error occurs during an update, auto-rollback is triggered to return the cloud to its pre-update state. After an update you can check for any functional impacts on the cloud. If everything is fine you can commit the update, which deletes the old containers and old images from the nodes. Should you see any functional cloud impact you can perform a manual rollback to start the old containers again.

Before you begin a container update, keep the following in mind:

- Updates are not supported for registry-related containers and authorized_keys.
- You cannot roll back the repo containers on the management node to an older version after they are updated because rollbacks will delete node packages and might cause the cloud to destabilize.
- To prevent double-faults, a cloud sanity check is performed before the update is started. A cloud sanity check is performed as the last step of the update.

The following table provides an overview to the methods to start the OpenStack update using Cisco VIM. The Internet options refer to management node connectivity to the Internet. If your management server lacks Internet access, you must have a staging server with Internet access to download the Cisco VIM installation artifacts to a USB stick. Cisco recommends selecting one method and staying with it for the full pod lifecycle.

Table 2: OpenStack Update Options

	Without Cisco VIM Insight	With Cisco VIM Insight
Without Internet	<ul style="list-style-type: none"> • Prepare the USB on a staging server • Plug the USB into the management node. • Follow the update steps in the update without Internet procedure. 	<ul style="list-style-type: none"> • Prepare the USB on a staging server • Plug the USB into the management node. • Follow the update steps in the update without Internet procedure.
With Internet	<ul style="list-style-type: none"> • Download the .tgz file from the registry. • Follow the update steps in the update with Internet procedure. 	<ul style="list-style-type: none"> • Download the .tgz file from the registry. • Follow the update steps in the update with Internet procedure.

Updating Cisco VIM Software Using a USB

The following procedure tells you how to load the Cisco VIM installation files onto a Cisco NFVI management node that does not have Internet access. Installation files include: buildnode-K9.iso, mercury-installer.tar.gz, nova-libvirt.tar, registry-2.3.1.tar.gz, and respective checksums..

Before you begin

This procedure requires a CentOS 7 staging server (VM, laptop, or UCS server) with a 64 GB USB 2.0 stick. The staging server must have Internet access (wired access is recommended) to download the Cisco VIM installation files, which you will load onto the USB stick. You then use the USB stick to load the installation files onto the management node. The installation files are around 24 GB in size, downloading them to the USB stick might take several hours, depending on the speed of your Internet connection, so plan accordingly. Before you begin, disable the CentOS sleep mode.

Step 1 On the staging server, use yum to install the following packages:

- PyYAML (yum install PyYAML)
- python-requests (yum install python-requests)

Step 2 Connect to the Cisco VIM software download site using a web browser and login credentials provided by your account representative and download the **getartifacts.py** script from external registry.

```
# download the new getartifacts.py file (see example below)
curl -o getartifacts.py
https://<username>:<password>@cvim-registry.com/mercury-releases/mercury-rhel7-osp8/releases/<1.0.1>/getartifacts.py
```

```

curl -o getartifacts.py-checksum.txt
https://<username>:<password>@cvm-registry.com/mercury-releases/mercury-rhel7-osp8/releases/1.0.1/getartifacts.py-checksum.txt

# calculate the checksum and verify that with one in getartifacts.py-checksum.txt
sha512sum getartifacts.py

# Change the permission of getartificats.py
chmod +x getartifacts.py

```

Step 3 Run the **getartifacts.py** script. The script formats the USB 2.0 stick and downloads the installation artifacts. You will need to provide the registry username and password, the tag ID, and the USB partition on the staging server. For example:

To identify the USB drive, execute the **lsblk** command before and after inserting the USB stick. (The command displays a list of available block devices.) The output delta will help find the USB drive location. Provide the entire drive path in the **-d** option, instead of any partition.

```
sudo ./ getartifacts.py -t <tag_id> -u <username> -p <password> -d </dev/sdc>
```

Note Do not remove the USB stick while the synchronization is under way.

Step 4 Verify the integrity of the downloaded artifacts and the container images:

```

# create a directory
sudo mkdir -p /mnt/Cisco

# /dev/sdc is the USB drive, same as supplied in get artifacts.py python script
sudo mount /dev/sdc1 /mnt/Cisco
cd /mnt/Cisco

# execute the verification script
./test-usb

# failures will be explicitly displayed on screen, sample success output below
# sample output of ./test-usb execution with 2.2.x release
[root@mgmtnode Cisco]# ./test-usb
INFO: Checking the integrity of this USB stick
INFO: Checking artifact buildnode-K9.iso
INFO: Checking artifact mercury-version.txt
INFO: Checking artifact registry-2.3.1.tar.gz
INFO: Checking artifact nova-libvirt-K9.tar.gz
INFO: Checking required layers:
INFO: 395 layer files passed checksum.
[root@mgmtnode Cisco]#

```

Step 5 To resolve download artifact failures, unmount the USB and run the **getartifacts** command again with the **--retry** option:

```
sudo ./getartifacts.py -t <tag_id> -u <username> -p <password> -d </dev/sdc> --retry
```

Step 6 Mount the USB and then run the **test-usb** command to validate all the files are downloaded:

```

# /dev/sdc is the USB drive, same as supplied in get artifacts.py python script
sudo mount /dev/sda1 /mnt/Cisco
cd /mnt/Cisco

# execute the verification script
./test-usb

# In case of failures the out of the above command will explicitly display the same on the screen

```

Step 7 After the synchronization finishes, unmount the USB stick:

```
sudo umount /mnt/Cisco
```

Step 8 After the synchronization finishes, remove the USB stick from the staging server then insert it into the management node.

Step 9 Complete the following steps to import the Cisco NFVI installation artifacts onto the management node:

a) Identify the USB on the management node:

```
blkid -L Cisco-VIM
```

b) Mount the USB device on the management node:

```
mount < /dev/sdc > /mnt/  
cd /tmp/
```

c) Extract the import_artifacts.py script:

```
tar --no-same-owner -xvzf /mnt/mercury-installer.tar.gz
```

d) Unmount the USB device:

```
umount /mnt/
```

e) Import the artifacts:

```
cd /tmp/installer-< xxxx >/tools/  
./import_artifacts.sh
```

f) Change directory and remove /tmp/installer-< xxxx >

```
cd /root/  
rm -fr /tmp/installer-< xxxx >
```

Step 10 Verify the image version and change ID for the software update.

```
cat /var/cisco/artifacts/mercury-version.txt
```

Step 11 Execute the update from the old working directory:

```
cd $old_workspace/installer;  
ciscovim update --file /var/cisco/artifacts/mercury-installer.tar.gz
```

After the update is complete, use the newly created directory from here onwards (unless a rollback is planned).

Step 12 Commit the update by running the following command:

```
ciscovim commit # from the new workspace
```

Step 13 To revert the update changes before entering the commit command, enter:

```
ciscovim rollback # and then use older workspace
```

Note Do not run any other Cisco VIM actions while the update is underway.

In Cisco VIM 2.2.12, if updates bring in Kernel changes, then the reboot of the compute node with VNFs in ACTIVE state is postponed. This is done to mitigate the unpredictability of data plane outage when compute nodes go for a reboot for the kernel changes to take effect, during the rolling upgrade process.

At the end of ciscovim update, the CVIM orchestrator displays the following message on the console and logs:

```
Compute nodes require reboot Kernel updated  
<compute_1_with_VM_running>  
<compute_3_with_VM_running>
```

```
<compute_4_with_VM_running>
<compute_12_with_VM_running>
```

After the Kernel update on Management node, reboot the compute node before proceeding. The logs for this run are available in `<mgmt._ip_address>:/var/log/mercury/<UUID>`

Note As the redundancy in controller, and storage nodes are built into the product, the reboot of those nodes are automatic during the software update. Also, computes that does not have any VNFs in ACTIVE state, gets automatically rebooted during software update

Updating Cisco VIM Software Using Network Installation

Step 1 From the download site that is provided by your Cisco account representative, download the mercury-installer.gz

```
curl -o mercury-installer.tar.gz
https://{username}:{password}@cvim-registry.cisco.com/
mercury-releases/mercury-rhel7-osp10/releases/{release number}/
mercury-installer.tar.gz
```

The link to the tar ball preceding is an example.

Step 2 Execute the update from the old working directory:

Note Do not run any other Cisco VIM actions while the update is underway.

```
cd /root/installer-<tagid>
ciscovim update -file /root/mercury-installer.tar.gz
```

After the update is complete, use the newly created directory from here onwards (unless a rollback is planned).

Step 3 Commit the update by running the following command:

```
ciscovim commit
```

Step 4 To revert the update changes before entering the commit command, enter:

```
ciscovim rollback # and then use older workspace
```

In Cisco VIM, if updates bring in Kernel changes, then the reboot of the compute node with VNFs in ACTIVE state is postponed. This is done to mitigate the unpredictability of data plane outage when compute nodes go for a reboot for the kernel changes to take effect, during the rolling upgrade process.

At the end of ciscovim update, the CVIM orchestrator displays the following message on the console and logs:

```
Compute nodes require reboot Kernel updated
<compute_1_with_VM_running>
<compute_3_with_VM_running>
<compute_4_with_VM_running>
<compute_12_with_VM_running>
```

After the Kernel update on the Management node, reboot the compute node before proceeding

The logs for this run are available in `<mgmt._ip_address>:/var/log/mercury/<UUID>`

Note The redundancy in controller, and storage nodes are built into the product, the reboot of those nodes are automatic during the software update. Also, computes that does not have any VNFs in ACTIVE state, gets automatically rebooted during the software update. To monitor and reboot the compute nodes through `ciscovim cli`, refer to the sections titled “Managing Reboot of Cisco VIM Nodes:” and “Managing Reboot Status of Cisco VIM Nodes”, in the later part of this guide. It should be noted no pod management operation is allowed till reboot of all CVIM nodes are successful.

Upgrading Containers in a Running Cisco VIM Cloud

Cisco VIM 2.2 allows you to upgrade all OpenStack services, infrastructure services such as RabbitMQ, MariaDB, HAProxy, and management node containers such as Cobbler, ELK, VMTP and repo containers. You can upgrade to new releases of OpenStack without redeploying the Cisco NFVI stack from the beginning. During upgrade, you can expect limited service impact as the upgrade is run serially on component by component (one node at a time).

Cisco VIM 2.2 supports upgrade from a known version of VIM running Liberty (1.0.43) to the current version of Newton (2.2.24). As OpenStack does not support the skipping of major releases during upgrade, the VIM upgrade internally moves the stack to Mitaka and then to Newton release of OpenStack.

As part of the VIM cloud upgrade,

- The `runner.py` script is used to automatically upgrade the REST API server managing the VIM orchestrator.
- The `setup_data.yaml` is automatically translated so that the `setup_data.yaml` file is compatible to the target release version.

Before you begin a container update, consider the following points:

- Plan for the downtime as the upgrade involves moving the Kernel version.
- Updates are not supported for registry-related containers and `authorized_keys`.
- The repo containers on the management node cannot be rolled back to an older version after the upgrade, as the rollbacks will delete the node packages and destabilize the cloud.
- A cloud sanity check is performed before the update is started to prevent double-faults. A cloud sanity check is performed as the last step of the update.

Before you begin a Pod upgrade, keep the following in mind:

- There is no roll-back in Upgrades, so it is better to stage it in the lab, and test it few times as issues related to specific customer environment might surface.
- The upgrade script, `vim_upgrade_orchestrator.py`, is available as part of the 2.2 artifacts and needs to be copied to `/root/` location before starting the execution.
- For disconnected upgrade, 2 USBs 2.0 (64GB) should be pre-populated with artifacts from 1.5.19 and 2.2.x.
- Upgrade from 1.0.43 to 2.2.x is restricted to specific starting and end point.
- Upgrade of the cloud is supported in both connected and disconnected mode.
- In 2.2.x, UCSD is no longer supported, instead the UI has been replaced by VIM Insight; Post Upgrade customer should bring up the Insight service on its own and register the pod to it

- We recommended you not to change the upgrade mode and the install mode.
- Upgrade is a one-way operation (there no rollback); so planning should be done before executing the upgrade. In the off chance, if one faces an issue, reach out to Cisco TAC/BU to recover the cloud.
- Prior to upgrade, ensure that the size of the storage network is greater than the total number of nodes (control, compute and ceph) the pod has. This is extremely important; failure to meet this criterion can cause the failure of the pre-upgrade check, and can prevent the upgrade from moving forward.

Following are the steps on how to get this conversion done before the upgrade commences.

- Make a copy of the golden `~/openstack-configs/setup_data.yaml` from the installer as a different file name For Example: `~/setup_data-reconfigure-storage-net.yaml`
- Edit the storage network segment in `~/setup_data-reconfigure-storage-net.yaml` copy to the desired subnet

```
NETWORKING:
.....
- gateway: 17.16.99.1
pool: [17.16.99.2 to 17.16.99.6]
segments: [storage]
subnet: 17.16.99.0/24
vlan_id: '3005'
```



Note Do not expand the same storage subnet use the different subnet like:

- `# cd ~/installer-1.0.43/tools`
- `./reconfigure_storage_network.sh --setup_file /root/setup_data-reconfigure-storage-net.yaml`

It takes 10 minutes to complete on a 3 controller, 3 compute, 3 storage node approximately.

At a high level, the upgrade script, `vim_upgrade_orchestrator.py`, is broken into three steps with logic to abort on fail. In case of failure, it is important to call Cisco support and not recover the cloud on your own.

The following are the three high level steps into which the `vim_upgrade_orchestrator.py` is broken into:

- **Pre-Upgrade Check**
 - Registry connectivity (if connected install).
 - Setup_data pre check: No UCSM_PLUGIN, sufficient storage pool size.
 - Backup setup_data.yaml, before performing translation.
 - Check and Update INSTALL_MODE in setup_data.yaml (connected or disconnected);
 - Check the storage network size.
 - run `cloud - sanity` from stable/liberty.
 - Check for reachability to all nodes including compute, controller and storage.
- **Upgrade to 1.5.18 (Mitaka):**
 - Delete UCSD instance if its running on the management node.

- Delete old version of nova-libvirt on the management node.
 - Auto-translation of setup_data for Mitaka.
 - Upgrade to Mitaka.
 - Check for reachability to all nodes (compute, controller and storage).
- **Upgrade to 2.2.0 (Newton).**
- Upgrade to Newton.
 - Backup of Management Node.
 - run cloud- sanity from stable/newton.
 - Check for reachability to all nodes (compute, controller and storage)

For a cloud running with VTS; additional steps need to be taken at this point, as VTC is not managed by Cisco VIM.

At a high level listed following are the manual steps:

- Power down the VTS 2.3 master and slave instances.
- Keep the VTS 2.3 XRNC/XRVR master and slave instance running.
- Bring up new VTS 2.5 master and slave instance.
- Login to master and slave VTS 2.5 UI and change admin password.
- Enable the VTS High availability with same assigned VIP IP address in VTS 2.3 deployment.
- Bring up new VTSR 2.5 master and slave instance.
- Login to VTS HA VIP IP address and execute the Underlay Loopback and OSPF Template.
- Power down the VTS 2.3 XRNC/XRVR Master and slave instance before performing vtc upgrade only.

After executing the manual steps, execute the `vim_upgrade_orchestrator.py` with `vtcupgradeonly` option to have the VIM cloud working with VTS 2.5.

- Connect to the CIMC of the management node and validate the boot-order list SDCARD as the first choice.
- Power-cycle the management node to complete the management node upgrade.
- Manually move the Pod from Software Raid to Hardware Raid.

The following table provides an overview to the methods to start the OpenStack update using Cisco VIM. The Internet options refer to management node connectivity to the Internet. If your management server lacks Internet access, you must have a staging server with Internet access to download the Cisco VIM installation artifacts to a USB stick. We recommend you to select one method and stay with it for the full pod lifecycle.

Upgrade Method	Without Cisco VIM Insight
Without Internet	<ul style="list-style-type: none"> • Prepare 2 USB 2.0 (64G) on a staging server and populate them with 1.5.x and 2.2.0 artifacts. • Plug both the USB into the management node. • Copy the vim_upgrade_orchestrator.py and follow the upgrade steps in the upgrade without Internet procedure.
With Internet	Copy the vim_upgrade_orchestrator.py and follow the upgrade steps in the upgrade without Internet procedure

Upgrading VIM Software Using a USB

The following procedure tells you how to load the Cisco VIM installation files onto a Cisco NFVI management node that does not have Internet access. Installation files include: build node-K9.iso, mercury-installer.tar.gz, nova-libvirt.tar, registry-2.3.1.tar.gz, and respective checksums.

Before you begin

This procedure requires a CentOS 7 staging server (VM, laptop, or UCS server) with two 64 GB USB 2.0 stick. The staging server must have Internet access(wired access is recommended) to download the Cisco VIM installation files, which you will load onto the USB stick. You then use the USB stick to load the installation files onto the management node. The installation files are around 24 GB in size, downloading them to the USB stick might take several hours, depending on the speed of your Internet connection, so plan accordingly. Before you start, disable the CentOS sleep mode.

Step 1 On the staging server, use yum to install the following packages:

- PyYAML (yum install PyYAML)
- python-requests (yum install python-requests)

Step 2 Connect to the Cisco VIM software download site using a web browser and log in to the username and the password that are provided by your account representative and download the **getartifacts.py** script from external registry.

```
# download the new getartifacts.py file (see example below) curl -o getartifacts.py
https://<username>:<password>@cvm-registry.com/mercury-releases/mercury-rhel7-osp9/releases/<1.5.2>/getartifacts.py

curl -o getartifacts.py-checksum.txt
https://<usemane>:<password>@cvm-registry.com/mercury-releases/mercury-rhel7-osp9/releases/1.5.2/getartifacts.py-checksum.txt

# calculate the checksum and verify that with one in getartifacts.py-checksum.txt sha512sum
getartifacts.py

# Change the permission of getartificats.py
chmod +x getartifacts.py
```

Step 3 Run the `getartifacts.py` script. The script formats the USB 2.0 stick and downloads the installation artifacts. You will need to provide the registry username and password, the tag ID, and the USB partition on the staging server. For example:

To identify the USB drive, execute the `lsblk` command before and after inserting the USB stick. (The command displays a list of available block devices.) The output data helps find the USB drive location. Provide the entire drive path in the `-d` option, instead of any partition.

```
sudo ./getartifacts.py -t <tag_id> -u <username> -p <password> -d </dev/sdc>
```

Note Do not remove the USB stick while the synchronization is going on.

Step 4 Verify the integrity of the downloaded artifacts and the container images:

```
# create a directory sudo mkdir -p /mnt/Cisco

# /dev/sdc is the USB drive, same as supplied in get artifacts.py python script sudo mount /dev/sdc1
/mnt/Cisco
cd /mnt/Cisco

# execute the verification script
./test-usb

# failures will be explicitly displayed on screen, sample success output below
# sample output of ./test-usb execution with 2.2.x release [root@mgmtnode Cisco]# ./test-usb
INFO: Checking the integrity of this USB stick INFO: Checking artifact buildnode-K9.iso
INFO: Checking artifact mercury-version.txt INFO: Checking artifact registry-2.3.1.tar.gz INFO:
Checking artifact nova-libvirt-K9.tar.gz INFO: Checking required layers:
INFO: 395 layer files passed checksum. [root@mgmtnode Cisco]#
```

Step 5 To resolve download artifact failures, unmount the USB and run the `getartifacts` command again with the `--retry` option:

```
sudo ./getartifacts.py -t <tag_id> -u <username> -p <password> -d </dev/sdc> --retry
```

Step 6 Mount the USB and then run the `test-usb` command to validate all the files are downloaded:

```
# /dev/sdc is the USB drive, same as supplied in get artifacts.py python script
sudo mount /dev/sda1 /mnt/Cisco
cd /mnt/Cisco

# execute the verification script
./test-usb

# In case of failures the out of the above command will explicitly display the same on the screen
```

Step 7 After the synchronization finishes, unmount the USB stick:

```
sudo umount /mnt/Cisco
```

Step 8 After the synchronization finishes, remove the USB stick from the staging server then insert it into the management node.

Step 9 Repeat step 2 to step 8 for pre-population of CVIM 2.2.x artifacts onto the second USB

Step 10 Insert the 2 pre-populated USBs into the management node of the pod running 1.0.41.

Step 11 Copy the `vim_upgrade_orchestrator.py` script available in CVIM 2.2 artifacts in the `/root/` folder of the management node of the pod running 1.0.41

Step 12 Execute the update from the `/root/` location:

```
# cd /root/
# ./vim_upgrade_orchestrator.py -i disconnected [-y] # -y if you don't want any interactive mode
```

After the upgrade is complete, use the newly created directory from here onwards.

Note Upgrade process takes several hours (> 6 hours), so execute this process in a VNC. Do not run any other Cisco VIM actions while the upgrade is underway.

Step 13 Copy the management node backup that is created during the upgrade, into a separate server through rsync (see chapter 12 for details).

Step 14 Check that the SDCARD is state as priority 1 for the boot order, from the CIMC of the management node. If not, set it accordingly. Reboot the management node, and wait for it to come all the way up.

Step 15 If VTS is running on the pod, manually transfer the VTC to 2.5

Step 16 Execute the update from the /root/ location:

```
# cd /root/
# ./vim_upgrade_orchestrator.py -i disconnected -s VTCSSHUSERNAME -p VTCSSHPASSWORD -vtsupgradeonly [-y]
```

Note "-y" is if you don't want any interactive mode

Step 17 Manually switch the management node of the pod to Hardware Raid. At a high level following steps need to be followed on the management node:

- Go to CIMC and set HDD as the top boot order.
- Enable the HBA on Bios setting from the CIMC. For procedure details, refer to Setting up the UCS C-Series Pod section in the 2.2 Install Guide.
- Initiate backup and restore of the management node with the current upgraded image version. For details on how to do it, refer to Chapter 12, where the backup and restore of the management node is listed.

Step 18 Move the pod to Hardware Raid, execute the following:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
```

```
Update the kickstart line in setup_data in /root/Save/:
kickstart:
  block_storage: storage-flexflash-c240m4.ks
  compute: compute-flexflash-c220m4.ks
  control: control-flexflash-c220m4.ks
to
kickstart:
  control: ucs-b-and-c-series.ks
  compute: ucs-b-and-c-series.ks
block_storage: storage-flexflash-c240m4.ks
```

```
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> run -perform 1
```

- Then for each compute node, do a remove, followed by add of the same node; (don't forget to power on the server after the remove); you can do remove followed by add of more than 1 compute node at a time.
- For every controller node, execute the replace controller, one at a time

Note To verify that the systems (compute and controllers) have moved to hardware RAID, execute the following in each server after the above steps

```
# /opt/MegaRAID/storcli/storcli64 /c0 /v0 show | grep Status
and verify that "Status = Success" is displayed.
```

Upgrading Cisco VIM Software Using Network Installation

Step 1 From the download site that is provided by your Cisco account representative, download the `vim_upgrade_orchestrator.py`
`curl -o vim_upgrade_orchestrator.py https://{username}:{password}@cvm-registry.cisco.com/mercury-releases/mercury-rhel7-osp10/releases/{release number}/vim_upgrade_orchestrator.py`. The link to the tar ball preceding is an example.

Step 2 Execute the upgrade from `/root/` directory:

```
$ cd /root/
$ ./vim_upgrade_orchestrator.py -i connected
```

Note Do not run any other Cisco VIM actions while the update is going on.

After the upgrades are complete, use the newly created folder.

VM Resizing

VM resize is the process of changing the flavor of an existing VM. Thus, using VM resize you can upscale a VM according to your needs. The size of a VM is indicated by the flavor based on which the VM is launched.

Resizing an instance means using a different flavor for the instance.

By default, the resizing process creates the newly sized instance on a new node, if more than one compute node exists and the resources are available. By default, the software, allows you to change the RAM size, VDISK size, or VCPU count of an OpenStack instance using **nova resize**. Simultaneous or individual adjustment of properties for the target VM is allowed. If there is no suitable flavor for the new properties of the VM, you can create a new one.

```
nova resize [--poll] <server> <flavor>
```

The resize process takes some time as the VM boots up with the new specifications. For example, the Deploying a Cisco CSR (size in MB) would take approximately 60mins. After the resize process, execute `nova resize-confirm <server>` to overwrite the old VM image with the new one. If you face any issue, you can revert to the old VM using the `nova-resize-revert <server>` command. At this point, you can access the VM through SSH and verify the correct image is configured.



Note The OpenStack **shutdown** the VM before the resize, so you have to plan for a **downtime**.



Note We recommend you not to resize a vdisk to a smaller value, as there is the risk of losing data.

Nova Migrate

The `nova migrate` command is used to move an instance from one compute host to another compute host. The scheduler chooses the destination compute host based on the availability of the zone settings. This process does not assume that the instance has shared storage available on the target host.

To initiate the cold migration of the VM, you can execute the following command:

```
nova migrate [--poll] <server>
```

The VM migration can take a while, as the VM boots up with the new specifications. After the VM migration process, you can execute `nova resize-confirm <server> --to` to overwrite the old VM image with the new one. If you encounter a problem, use the `nova-resize-revert <server>` command to revert to the old VM image. At this point, access the VM through SSH and verify the correct image is configured.



Note The OpenStack **shutdown** the VM before the migrate, so plan for a **downtime**.



CHAPTER 2

Cisco VIM REST API

The following topics explain how to use the Cisco VIM REST API to manage Cisco NFVI.

- [Overview to Cisco VIM REST API, on page 47](#)
- [Cisco VIM REST API Resources, on page 48](#)

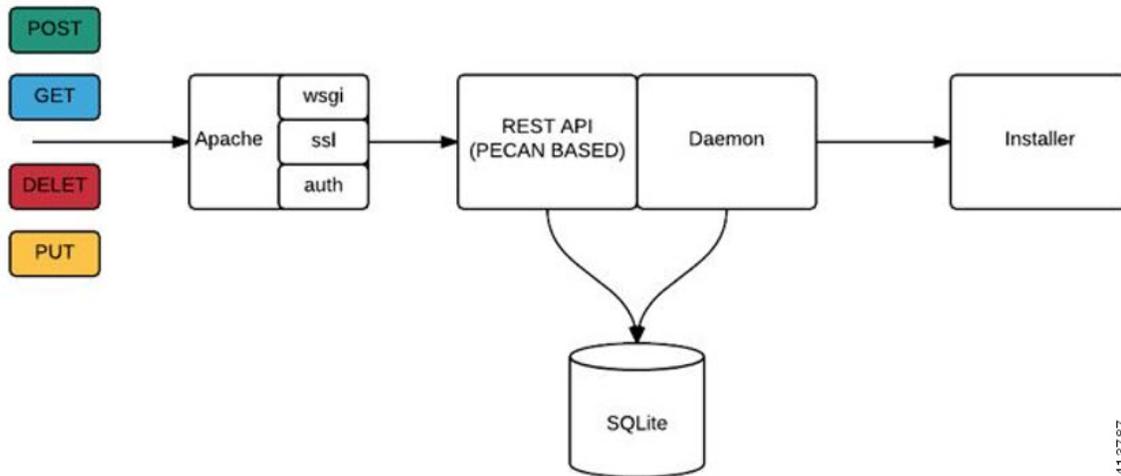
Overview to Cisco VIM REST API

Cisco VIM provides a Representational State Transfer (REST) API that is used to install, expand, and update Cisco VIM. Actions performed using the REST APIs are:

- Install Cisco VIM on Cisco NFVI pods
- Add and delete pods to and from Cisco NFVI installations
- Update Cisco VIM software
- Replace controller nodes
- Perform cloud maintenance operations
- Run cloud validations using Virtual Machine ThroughPut (VMTP), a data path performance measurement tool for OpenStack clouds

The following figure shows the Cisco VIM REST API flow.

Figure 7: Cisco VIM REST API Flow



413787

The Cisco VIM REST API security is provided by the Secure Sockets Layer (SSL) included on the Apache webserver. The Pecan-based web application is called by `mod_wsgi`, which runs the Rest API server. The Pecan REST API server requires a username and password to authorize the REST API server requests. Apache handles the authorization process, which authorizes the request to access the Pecan web application. Use the Cisco VIM API to upload a new `setup_data.yaml` file, and start, stop, and query the state of the installation. You can use it to manage the cloud, add and remove compute and Ceph nodes, and replace the controller nodes. A REST API to launch VMTP (L2/L3 data plane testing) and CloudPulse is also provided.

The Cisco VIM REST API is enabled by default in the management node, if you are using the supplied Cisco VIM `buildnode.iso`. You can access API server on the `br_api` interface on port 8445. Authentication is enabled by default in the web service.

The API endpoints can be reached with the following URL format:

```
https://<Management_node_api_ip>:8445
```

The API endpoint expects a basic authentication which is enabled by default in the management node. The authentication credentials are found in `/opt/cisco/ui_config.json` in the management node.

The following contents show the Sample `ui_config.json`:

```
{
  "Kibana-Url": "http://10.10.10.10:5601",
  "RestAPI-Url": "https:// 10.10.10.10:8445",
  "RestAPI-Username": "admin",
  "RestAPI-Password": "a96e86ccb28d92ceb1df",
  "RestDB-Password": "e32de2263336446e0f57",
  "BuildNodeIP": "10.10.10.10"
}
```

Cisco VIM REST API Resources

Setupdata

REST wrapper for setupdata. Provides methods for listing, creating, modifying and deleting setupdata.

Retrieving the setupdata

Resource URI

Verb	URI
GET	/v1/setupdata

Example

JSON Request

```
GET /v1/setupdata
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{"setupdatas": [{
  "status": "Active",
  "name": "GG34",
  "uuid": "123"
  "meta": {
    "user": "root"
  },
  "jsondata": {
    .....
  }
}]}
```

Creating the setupdata

Resource URI

Verb	URI
POST	/v1/setupdata

Example

JSON Request

```
POST /v1/setupdata
Accept: application/json
```

```
{
  "name": "GG34",
  "uuid": "123"
  "meta": {
    "user": "root"
  },
  "jsondata": {
    .....
  }
}
```

JSON Response

```
201 OK
Content-Type: application/json
{
  "status": "Active",
```

```

        "name": "GG34",
        "uuid": "123"
        "meta": {
            "user": "root"
        },
        "jsondata": {
            .....
        }
    }
}

400 Bad Request
Content-Type: application/json
{
    "debuginfo": null
    "faultcode": "Client"
    "faultstring": "Error"
}

409 CONFLICT
Content-Type: application/json
{
    "debuginfo": null
    "faultcode": "Client"
    "faultstring": "Error"
}

```

Retrieving a single setupdata

Resource URI

Verb	URI
GET	/v1/setupdata/(id)

Property:

id - the id of the setupdata to be queried.

Example

JSON Request

```

GET /v1/setupdata/123
Accept: application/json

```

JSON Response

```

200 OK
Content-Type: application/json
{
    "status": "Active",
    "name": "GG34",
    "uuid": "123"
    "meta": {
        "user": "root"
    },
    "jsondata": {
        .....
    }
}

404 NOT FOUND
Content-Type: application/json

```

```
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Setupdata could not be found."
}
```

Updating a setupdata

Resource URI

Verb	URI
PUT	/v1/setupdata/(id)

Property:

id - the id of the setupdata to be updated.

Example

JSON Request

```
PUT /v1/setupdata/123
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "status": "Active",
  "name": "GG34",
  "uuid": "123"
  "meta": {
    "user": "root"
  },
  "jsondata": {
    .....
  }
}
```

```
404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Setupdata could not be found."
}
```

Deleting a setupdata

Resource URI

Verb	URI
DELETE	/v1/setupdata/(id)

Property:

id - the id of the setupdata to be deleted.

Example

JSON Request

```
DELETE /v1/setupdata/123
Accept: application/json
```

JSON Response

```
204 NO CONTENT
Returned on success
```

```
404 NOT FOUND
Content-Type: application/json
```

```
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Setupdata could not be found."
}
```

```
400 BAD REQUEST
Content-Type: application/json
```

```
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Setupdata cannot be deleted when it is being used by an installation"
}
```

Install resource

REST wrapper for install. Provides methods for starting, stopping, and viewing the status of the installation process.

Return a list of installation

Resource URI

Verb	URI
GET	/v1/install

Example

JSON Request

```
GET /v1/install
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
```

```
{"installs": [{
  "ceph": "Skipped",
  "uuid": "123",
  "setupdata": "345",
  "vmtpresult": "{
    "status": "PASS",
    "EXT_NET": []
  }",
  "baremetal": "Success",
```

```

        "orchestration": "Success",
        "validationstatus": "{
            "status": "PASS",
            "Software_Validation": [],
            "Hardware_Validation": []
        }",
        "currentstatus": "Completed",
        "validation": "Success",
        "hostsetup": "Success",
        "vmtp": "Skipped"
    }
}

```

Create an installation

Resource URI

Verb	URI
POST	/v1/install

Example

JSON Request

```

GET /v1/install
Accept: application/js
{
    "setupdata": "123",
    "stages": [
        "validation",
        "bootstrap",
        "runtimevalidation",
        "baremetal",
        "orchestration",
        "hostsetup",
        "ceph",
        "vmtp"
    ]
}

```

JSON Response

```

201 CREATED
Content-Type: application/json
{
    "ceph": "Skipped",
    "uuid": "123",
    "setupdata": "345",
    "vmtpresult": "{
        "status": "PASS",
        "EXT_NET": []
    }",
    "baremetal": "Success",
    "orchestration": "Success",
    "validationstatus": "{
        "status": "PASS",
        "Software_Validation": [],
        "Hardware_Validation": []
    }",
    "currentstatus": "Completed",
    "validation": "Success",
}

```

```

    "hostsetup": "Success",
    "vmtp": "Skipped"
  }

409 CONFLICT
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Install already exists"
}

```

Retrieve the installation

Resource URI

Verb	URI
GET	/v1/install/{id}

Property:

id - the id of the install to be queried.

Example

JSON Request

```

GET /v1/install/345
Accept: application/js

```

JSON Response

```

200 OK
Content-Type: application/json
{
  "ceph": "Skipped",
  "uuid": "123",
  "setupdata": "345",
  "vmtpresult": "{
    \"status\": \"PASS\",
    \"EXT_NET\": []
  }",
  "baremetal": "Success",
  "orchestration": "Success",
  "validationstatus": "{
    \"status\": \"PASS\",
    \"Software_Validation\": [],
    \"Hardware_Validation\": []
  }",
  "currentstatus": "Completed",
  "validation": "Success",
  "hostsetup": "Success",
  "vmtp": "Skipped"
}

404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
}

```

```

    "faultcode": "Client"
    "faultstring": "Install doesn't exists"
  }

```

Stop the installation

Resource URI

Verb	URI
DELETE	/v1/install/{id}

Property:

id - the id of the install to be stopped.

Example

JSON Request

```

DELETE /v1/install/345
Accept: application/js

```

JSON Response

```

204 NO CONTENT
Content-Type: application/json

404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Install doesn't exists"
}

```

Nodes

Getting a list of nodes

Resource URI

Verb	URI
GET	/v1/nodes

Example

JSON Request

```

Get /v1/nodes
Accept: application/js

```

JSON Response

```

200 OK
Content-Type: application/json
{
  "nodes": [
    [

```

```

    "status": "Active",
    "uuid": "456",
    "setupdata": "123",
    "node_data": "{
      "rack_info": {
        "rack_id": "RackA"
      },
      "cimc_info": {
        "cimc_ip": "10.10.10.10"
      },
      "management_ip": "7.7.7.10"
    }",
    "updated_at": null,
    "mtype": "compute",
    "install": "345",
    "install_logs": "logurl",
    "created_at": "2016-0710T06:17:03.761152",
    "name": " compute-1"
  }
}

```

Add new nodes

The nodes are in compute or block_storage type. Before adding the nodes to the system, the name of the nodes and other necessary information like cimc_ip and rackid must be updated in the setupdata object. If the setupdata object is not updated, the post call will not allow you to add the node.

Resource URI

Verb	URI
POST	/v1/nodes

Example

JSON Request

```

POST /v1/nodes
Accept: application/js
{
  "name" : "compute-5"
}

```

JSON Response

```

201 CREATED
Content-Type: application/json
{
  "status": "ToAdd",
  "uuid": "456",
  "setupdata": "123",
  "node_data": "{
    "rack_info": {
      "rack_id": "RackA"
    },
    "cimc_info": {
      "cimc_ip": "10.10.10.10"
    },
    "management_ip": "7.7.7.10"
  }",
}

```

```

      "updated_at": null,
      "mtype": "compute",
      "install": "345",
      "install_logs": "logurl",
      "created_at": "2016-0710T06:17:03.761152",
      "name": " compute-1"
    }
  }

```

Retrieve information about a particular node

Resource URI

Verb	URI
GET	/v1/nodes{id}

Property:

id - the id of the node to be queried.

Example

JSON Request

```

POST /v1/nodes
Accept: application/js

```

JSON Response

```

200 OK
Content-Type: application/json
{
  "status": "Active",
  "uuid": "456",
  "setupdata": "123",
  "node_data": "{
    "rack_info": {
      "rack_id": "RackA"
    },
    "cimc_info": {
      "cimc_ip": "10.10.10.10"
    },
    "management_ip": "7.7.7.10"
  }",
  "updated_at": null,
  "mtype": "compute",
  "install": "345",
  "install_logs": "logurl",
  "created_at": "2016-0710T06:17:03.761152",
  "name": " compute-1"
}

404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Node doesn't exists"
}

```

Remove a node

The node that must be deleted must be removed from the setupdata object. Once the setupdata object is updated, you can safely delete of the node. The node object will not be deleted until it calls the remove node backend and succeeds.

Resource URI

Verb	URI
DELETE	/v1/nodes{id}

Property:

id - the id of the node to be removed.

Example

JSON Request

```
DELETE /v1/nodes/456
Accept: application/js
```

JSON Response

```
204 ACCEPTED
Content-Type: application/json

404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Node doesn't exists"
}
```

For clearing the database and deleting the entries in the nodes, the delete api is called with special parameters that are passed along with the delete request. The JSON parameters are in the following format.

JSON Request

```
DELETE /v1/nodes/456
Accept: application/js
{
  "clear_db_entry": "True"
}
```

JSON Response

```
204 ACCEPTED
Content-Type: application/json

404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Node doesn't exists"
}
```



Note This is done only if the node is deleted from the REST API database. The failure reason of the node must be rectified manually apart from the API. True is a string and not a boolean in the above line.

Replace a controller

Resource URI

Verb	URI
PUT	/v1/nodes{id}

Property:

id - the id of the controller to be replaced.

Example

JSON Request

```
PUT /v1/nodes/456
Accept: application/js
```

JSON Response

```
200 OK
Content-Type: application/json

404 NOT FOUND
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Node doesn't exists"
}
```

Offline validation

REST wrapper does the offline validation of setupdata. This will only do S/W Validation of the input setupdata.

Create an offline validation operation

Resource URI

Verb	URI
POST	/v1/offlinevalidation

Example

JSON Request

```
POST /v1/offlinevalidation
Accept: application/json
{
  "jsondata": "... .."
}
```

JSON Response

```

201 CREATED
Content-Type: application/json
{
  "status": "NotValidated",
  "uuid": "bb42e4ba-c8b7-4a5c-98b3-1f384aae2b69",
  "created_at": "2016-02-03T02:05:28.384274",
  "updated_at": "2016-02-03T02:05:51.880785",
  "jsondata": "{}",
  "validationstatus": {
    "status": "PASS",
    "Software_Validation": [],
    "Hardware_Validation": []
  }
}

```

Retrieve the results of offline validation

Resource URI

Verb	URI
GET	/v1/offlinevalidation

Property:

id - the id of the node to be queried.

Example

JSON Request

```

GET /v1/offlinevalidation/789
Accept: application/json

```

JSON Response

```

200 OK
Content-Type: application/json
{
  "status": " ValidationSuccess",
  "uuid": "bb42e4ba-c8b7-4a5c-98b3-1f384aae2b69",
  "created_at": "2016-02-03T02:05:28.384274",
  "updated_at": "2016-02-03T02:05:51.880785",
  "jsondata": "{}",
  "validationstatus": {
    "status": "PASS",
    "Software_Validation": [],
    "Hardware_Validation": []
  }
}

```

Update

Start an update process

Resource URI

Verb	URI
------	-----

POST	/v1/update
------	------------

Parameters:

- fileupload - "tar file to upload"
- filename - "Filename being uploaded"

Example

JSON Request

```
curl -sS -X POST --form
"fileupload=@Test/installer.good.tgz" --form
"filename=installer.good.tgz"
https://10.10.10.8445/v1/update
```



Note This curl request is done as a form request.

JSON Response

```
200 OK
Content-Type: application/json
{
  "update_logs": "logurl",
  "update_status": "UpdateSuccess",
  "update_filename": "installer-4579.tgz",
  "created_at": "2016-07-10T18:33:52.698656",
  "updated_at": "2016-07-10T18:54:56.885083"
}

409 CONFLICT
Content-Type: application/json
{
  "debuginfo": null
  "faultcode": "Client"
  "faultstring": "Uploaded file is not in tar format"
}
```

Rollback an update

Resource URI

Verb	URI
PUT	/v1/update

Example

JSON Request

```
PUT /v1/update
Accept: application/json
{
  "action": "rollback"
}
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "update_logs": "logurl",
  "update_status": "ToRollback",
  "update_filename": "installer-4579.tgz",
  "created_at": "2016-07-10T18:33:52.698656",
  "updated_at": "2016-07-10T18:54:56.885083"
}
```

Commit an update

Resource URI

Verb	URI
PUT	/v1/update

Example

JSON Request

```
PUT /v1/update
Accept: application/json
{
  "action": "commit"
}
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "update_logs": "logurl",
  "update_status": "ToCommit",
  "update_filename": "installer-4579.tgz",
  "created_at": "2016-07-10T18:33:52.698656",
  "updated_at": "2016-07-10T18:54:56.885083"
}
```

Retrieve the details of an update

Resource URI

Verb	URI
GET	/v1/update

Example

JSON Request

```
GET /v1/update
Accept: application/json
```

JSON Response

```

200 OK
Content-Type: application/json
{
  "update_logs": "logurl",
  "update_status": "UpdateSuccess",
  "update_filename": "installer-4579.tgz",
  "created_at": "2016-07-10T18:33:52.698656",
  "updated_at": "2016-07-10T18:54:56.885083"
}

```

Secrets

Retrieve the list of secrets associated with the OpenStack Setup

You can retrieve the set of secret password associated with the OpenStack setup using the above api. This gives the list of secrets for each service in OpenStack.

Resource URI

Verb	URI
GET	/v1/secrets

Example

JSON Request

```

GET /v1/secrets
Accept: application/json

```

JSON Response

```

200 OK
Content-Type: application/json
{
  "HEAT_KEYSTONE_PASSWORD": "xxxxx",
  "CINDER_KEYSTONE_PASSWORD": "xxxxxx",
  ...
  "RABBITMQ_PASSWORD": "xxxxxx"
}

```

OpenStack Configs

Retrieve the list of OpenStack configs associated with the OpenStack Setup

You can retrieve the set of OpenStack configs associated with the OpenStack setup using the above api. This gives the current settings of different configs like verbose logging, debug logging for different OpenStack services.

Resource URI

Verb	URI
GET	/v1/openstack_config

Example

JSON Request

```

GET /v1/openstack_config
Accept: application/json

```

JSON Response

```
200 OK
Content-Type: application/json
{
  "CINDER_DEBUG_LOGGING": false,
  "KEYSTONE_DEBUG_LOGGING": false,
  ...
  ...
  "NOVA_VERBOSE_LOGGING": true
}
```

Version

Retrieve the version of the Cisco Virtualized Infrastructure Manager.

Resource URI

Verb	URI
GET	/v1/version

Example

JSON Request

```
GET /v1/version
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{"version": "1.9.1"}
```

Health of the Management Node

Retrieve the health of the Management node

This api can be used to retrieve the health of the management node. It checks various parameters like partitions, space and so on.

Resource URI

Verb	URI
GET	/v1/health

Example

JSON Request

```
GET /v1/health
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "status": "PASS",
  "BuildNode Validation": {
    "Check Docker Pool Settings": {"status": "Pass", "reason": "None"}
    ...
    ...
  }
}
```

```
    }
  }
```

Hardware Information

REST wrapper to do hardware information of setupdata. This will return the hardware information of all hardware available in the setupdata.

Create a HWinfo operation

Resource URI

Verb	URI
GET	/v1/hwinfo

Example

JSON Request

```
POST /v1/hwinfo
Accept: application/json
{
  "setupdata": "c94d7973-2fcc-4cd1-832d-453d66e6b3bf"
}
```

JSON Response

```
201 CREATED
Content-Type: application/json
{
  "status": "hwinfoscheduled",
  "uuid": "928216dd-9828-407b-9739-8a7162bd0676",
  "setupdata": "c94d7973-2fcc-4cd1-832d-453d66e6b3bf",
  "created_at": "2017-03-19T13:41:25.488524",
  "updated_at": null,
  "hwinforeresult": ""
}
```

Retrieve the results of Hwinfo Operation

Resource URI

Verb	URI
GET	/v1/hwinfo/{id}

Property:

id - the id of the node to be queried.

Example

JSON Request

```
GET /v1/hwinfo/789
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "status": "hwinfosuccess",
  "uuid": "928216dd-9828-407b-9739-8a7162bd0676",
}
```

```

"setupdata": "c94d7973-2fcc-4cd1-832d-453d66e6b3bf",
"created_at": "2017-03-19T13:41:25.488524",
"updated_at": "2017-03-19T13:42:05.087491",
"hwinforeresult": {"172.29.172.73": {"firmware": .....
.....
.....
}

```

Release mapping Information

This api is used to see the list of Features included and list of options which can be reconfigured in the Openstack Setup.

Retrieve the release mapping information

Resource URI

Verb	URI
GET	/v1/releasemapping

JSON Request

```

GET /v1/releasemapping
Accept: application/json

```

JSON Response

```

200 OK
Content-Type: application/json
[
  {
    "SWIFTSTACK": {
      "feature_status": true,
    },
    "desc": "swift stack feature"
  }
  ,.....
  .....
]

```

POST Install operations

The following are the post install operations that can be carried on once the OpenStack installation is carried out successfully. It uses a common api. So only one operation is given as an example below:

1. reconfigure,
2. reconfigure -regenerate passwords
3. reconfigure -setpasswords,setopenstack_configs,
4. check-fernet-keys
5. period-rotate-fernet-keys
6. resync-fernet-keys
7. rotate-fernet-keys

Create a post install operation

Resource URI

Verb	URI
POST	/v1/misc

Example

JSON Request

```
POST /v1/misc
Accept: application/json
{"action": {"reconfigure": true}}
```

JSON Response

```
201 CREATED
Content-Type: application/json
{
  "uuid": "7e30a671-bacf-4e3b-9a8f-5a1fd8a46733",
  "created_at": "2017-03-19T14:03:39.723914",
  "updated_at": null,
  "operation_status": "OperationScheduled",
  "operation_logs": "",
  "operation_name": "{\"reconfigure\": true}"
}
```

Retrieve a status of the post install operation

Resource URI

Verb	URI
GET	/v1/misc

Example

JSON Request

```
GET /v1/misc
Accept: application/json
```

JSON Response

```
201 CREATED
Content-Type: application/json
{
  "uuid": "7e30a671-bacf-4e3b-9a8f-5a1fd8a46733",
  "created_at": "2017-03-19T14:03:39.723914",
  "updated_at": "2017-03-19T14:03:42.181180",
  "operation_status": "OperationRunning",
  "operation_logs": "xxxxxxxxxxxxxxxxxxxx",
  "operation_name": "{\"reconfigure\": true}"
}
```

In VIM 2.2, additional Rest APIs are introduced to support NFVBench, query hardware information and to get a list of optional and mandatory features that the pod supports.

Listed below are the details of the API.

NFVBench Network Performance Testing

Create NFVBench Run

Starts network performance test with provided configuration.

REST API To Create Fixed Rate Test

Verb	URI
Post	v1/nfvbench/create_ndr_pdr_test

Example

JSON Request

```
POST Request URL
/v1/nfvbench/create_fixed_rate_test
JSON Request:
{"nfvbench_request":
{
  "duration_sec": 20,
  "traffic_profile": [
    {
      "name": "custom",
      "l2frame_size": [
        "64",
        "IMIX",
        "1518"
      ]
    }
  ],
  "traffic": {
    "bidirectional": true,
    "profile": "custom"
  },
  "flow_count": 1000
}
}
```

JSON Response

```
201 CREATED
Content-Type: application/json
{
  "status": "not_run",
  "nfvbench_request":
  \{
    "duration_sec": 20,
    "traffic_profile": [
      {
        "name": "custom",
        "l2frame_size": [
          "64",
          "IMIX",
          "1518"
        ]
      }
    ]
  },
  "traffic": {
    "bidirectional": true,
    "profile": "custom"
  },
  "flow_count": 1000
}',
  "created_at": "2017-08-16T06:14:54.219106",
  "updated_at": null,
  "nfvbench_result": "",
  "test_name": "Fixed_Rate_Test"
}
```

Status Polling

Polling of NFVbench run status which is one of `nfvbench_running`, `nfvbench_failed`, `nfvbench_completed`.

Resource URI

Verb	URI
GET	v1/nfvbench/<test_name>

REST API To Get Fixed Rate Test Result

```
GET Request URL
/v1/upgrade/get_fixed_rate_test_result
JSON Response:
Check If NFVbench Test is running
200 OK
Content-Type: application/json
{
  "status": "nfvbench_running",
  "nfvbench_request": '{"traffic": {"bidirectional": true, "profile": "custom"},
"rate": "1000000pps",
"traffic_profile": [{"l2frame_size": ["1518"], "name": "custom"}], "duration_sec": 60,
"flow_count": 1000}',
  "nfvbench_result": ""
  "created_at": "2017-05-30T21:40:40.394274",
  "updated_at": "2017-05-30T21:40:41.367279",
}
```

```
Check If NFVbench Test is completed
200 OK
Content-Type: application/json
{
  "status": "nfvbench_completed",
  "nfvbench_request": '{"traffic": {"bidirectional": true, "profile": "custom"},
"rate": "1000000pps",
"traffic_profile": [{"l2frame_size": ["1518"], "name": "custom"}], "duration_sec": 60,
"flow_count": 1000}',
  "nfvbench_result": '{"status": "PROCESSED", "message": {"date": "2017-08-15 23:15:04"},
"nfvbench_version": "0.9.3.dev2", ...}'
  "created_at": "2017-05-30T21:40:40.394274",
  "updated_at": "2017-05-30T22:29:56.970779",
}
```

REST API to create NDR/PDR Test

```
POST Request URL
/v1/nfvbench/create_ndr_pdr_test
```

```
Accept: application/json
{"nfvbench_request":
{
  "duration_sec": 20,
  "traffic_profile": [
    {
      "name": "custom",
      "l2frame_size": [
        "64",
        "IMIX",
        "1518"
      ]
    }
  ],
  "traffic": {
```

```

        "bidirectional": true,
        "profile": "custom"
    },
    "flow_count": 1000
}
}

```

JSON Response

201 CREATED

Content-Type: application/json

```

{
  "status": "not_run",
  "nfvbench_request":
  \{
    "duration_sec": 20,
    "traffic_profile": [
      {
        "name": "custom",
        "l2frame_size": [
          "64",
          "IMIX",
          "1518"
        ]
      }
    ],
    "traffic": {
      "bidirectional": true,
      "profile": "custom"
    },
    "flow_count": 1000
  },
  "created_at": "2017-08-16T07:18:41.652891",
  "updated_at": null,
  "nfvbench_result": "",
  "test_name": "NDR_PDR_Test"
}

```

REST API To Get NDR/PDR Test Results

GET Request URL

/v1/ nfvbench/get_ndr_pdr_test_result

JSON Response:

If NFVbench NDR/PDR test is running

200 OK

Content-Type: application/json

```

{
  "status": "nfvbench_running",
  "nfvbench_request": '{"duration_sec": 20,
  "traffic": {"bidirectional": true, "profile": "custom"},
  "traffic_profile": [{"l2frame_size": ["64", "IMIX", "1518"], "name": "custom"}],
  "flow_count": 1000}',
  "nfvbench_result": ""
  "created_at": "2017-08-16T07:18:41.652891",
  "updated_at": "2017-09-30T22:29:56.970779",
}

```

If NFVbench NDR/PDR test is completed

200 OK

Content-Type: application/json

```

{
  "status": "nfvbench_completed",
  "nfvbench_request": '{"duration_sec": 20,

```

```

"traffic": {"bidirectional": true, "profile": "custom"},
"traffic_profile": [{"l2frame_size": ["64", "IMIX", "1518"], "name": "custom"}, {"flow_count":
  1000}],
  "nfvbench_result": '{"status": "PROCESSED",...}'
"created_at": "2017-08-16T07:18:41.652891",
"updated_at": "2017-09-30T22:29:56.970779",
}

```

REST API to Get Node Hardware Information

Rest API helps you to get the hardware information of all the nodes in the POD through CIMC/UCSM.

- Total Memory
- Firmware Info (Model, Serial Number)
- CIMC IP

GET Request URL

/v1/hwinfo

Output Response

```

{
  "hwinforeresult": [{"control-server-2": {"memory": {"total_memory": "131072"},
    "firmware": {"serial_number": "FCH1905V16Q", "fw_model": "UCSC-C220-M4S"},
    "cimc_ip": "172.31.230.100", "storage": {"num_storage": 4},
    "cisco_vic_adapters": {"product_name": "UCS VIC 1225"},
    "cpu": {"number_of_cores": "24"}, "power_supply": {"power_state": "on"}}
  ...
}

```

REST API to Get Mandatory Features Mapping

POST Request URL

/v1/releasemapping/mandatory_features_mapping

JSON Response:

```

{
  "mandatory": {
    "networkType": {
      "C": {
        "feature_status": true,
        "values": [{"name": "VXLAN/Linux Bridge", "value": "VXLAN/Linux Bridge"}],
        "insight_label": "Tenant Network",
        "desc": "Tenant Network"
      },
      "B": {
        "feature_status": true,
        "values": [{"name": "VXLAN/Linux Bridge", "value": "VXLAN/Linux Bridge"}],
        "insight_label": "Tenant Network",
        "desc": "Tenant Network"
      }
    },
    "cephMode": {
      "all": {
        "feature_status": true,
        "values": [{"name": "Central", "value": "Central"}],
        "insight_label": "Ceph Mode",
        "desc": "Ceph Mode"
      }
    },
    "podType": {
      "C": {

```

```

        "feature_status": true,
        "values": [{"name": "Fullon", "value": "fullon"}],
        "insight_label": "POD Type",
        "desc": "POD Type"
    },
    "B": {
        "feature_status": true,
        "values": [{"name": "Fullon", "value": "fullon"}],
        "insight_label": "POD Type",
        "desc": "POD Type"
    }
},
"installMode": {
    "all": {
        "feature_status": true,
        "values": [{"name": "Connected", "value": "connected"}, ],
        "insight_label": "Install Mode",
        "desc": "Install Mode"
    }
}
},
"platformType": [{"name": "B-series", "value": "B"}, {"name": "C-series", "value":
"C"}],
"postinstalllinks": {
    "view_cloudpulse": {"alwayson": true, "feature_status": true, "platformtype": "all",
"insight_label": "Run VMTP", "desc": "Cloudpluse"},
    "password_reconfigure": {"alwayson": true, "feature_status": true, "platformtype":
"all", "insight_label": "Reconfigure Passwords", "desc": "Reconfigure Passwords"}
}
}

```

REST API to Get Optional Features Mapping

POST Request URL
/v1/releasemapping/optional_features_mapping

JSON Response:

```

[
  {
    "SWIFTSTACK": {
      "feature_status": true,
      "insight_label": "Swiftstack",
      "repeated_redeployment": true,
      "reconfigurable": ["cluster_api_endpoint", "reseller_prefix", "admin_password",
"protocol"],
      "desc": "swift stack feature"
    }
  },
  {
    "heat": {
      "feature_status": true,
      "insight_label": "Heat",
      "repeated_redeployment": false,
      "reconfigurable": ["all"],
      "desc": "Openstack HEAT service"
    }
  },
  ... other features
]

```

Disk Maintenance information

REST wrapper to query information about RAID disks on Pod nodes. This will return the RAID disk information of all or a selection of RAID disks available in the Pod.

The disk management extension to the VIM REST API enables support for Disk Management actions

End Point	Type	Valid Args	Valid Arg Values	Example
/diskmgmt	GET	None	None	/v1/diskmgmt/
/diskmgmt/ check_disks	GET	None, args	All, control, compute	/v1/diskmgmt/check_disks/ /v1/diskmgmt/check_disks/? args=control,compute
/diskmgmt/ replace_disks	GET	None, args	All, control, compute	/v1/diskmgmt/replace_disks/ /v1/diskmgmt/replace_disks/? args=control,compute
/diskmgmt/server	GET	server_list, action	List of valid server names, check_disks, replace_disks	/v1/diskmgmt/server/?server_list= srv1,srv2&action=check_disks /v1/diskmgmt/server/?server_list= srv1&action=replace_disks

Get a Check disk operation

Resource URI

Verb	URI
GET	/v1/diskmgmt

Example

JSON Request

```
GET /v1/diskmgmt Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "add_as_spare_disks_results_list": [],
  "bad_disks_results_list": [],
  "fcfg_disks_results_list": [],
  "raid_results_list": [
    {
      "Num PDs": 4,
      "Num VDs": 1,
      "RAID health": "Opt",
      "RAID level": "RAID10",
      "RAID type": "HW",
      "VD health": "Opt1",
      "host": "EMC-Testbed",
      "role": "management",
      "server": "localhost"
    },
    {
      "Num PDs": 4,
      "Num VDs": 1,
      "RAID health": "Opt",
      "RAID level": "RAID10",
```

```

        "RAID type": "HW",
        "VD health": "Opt1",
        "host": "i13-20",
        "role": "control",
        "server": "15.0.0.7"
    },
    {
        "Num PDs": 4,
        "Num VDs": 1,
        "RAID health": "Opt",
        "RAID level": "RAID10",
        "RAID type": "HW",
        "VD health": "Opt1",
        "host": "i13-21",
        "role": "control",
        "server": "15.0.0.8"
    },
    {
        "Num PDs": 4,
        "Num VDs": 1,
        "RAID health": "Opt",
        "RAID level": "RAID10",
        "RAID type": "HW",
        "VD health": "Opt1",
        "host": "i13-22",
        "role": "control",
        "server": "15.0.0.5"
    },
    {
        "Num PDs": 4,
        "Num VDs": 1,
        "RAID health": "Opt",
        "RAID level": "RAID10",
        "RAID type": "HW",
        "VD health": "Opt1",
        "host": "i13-23",
        "role": "compute",
        "server": "15.0.0.6"
    },
    {
        "Num PDs": 4,
        "Num VDs": 1,
        "RAID health": "Opt",
        "RAID level": "RAID10",
        "RAID type": "HW",
        "VD health": "Opt1",
        "host": "i13-24",
        "role": "compute",
        "server": "15.0.0.10"
    }
],
"rbld_disks_results_list": [],
"spare_disks_results_list": []
}

```

Get a Check disk operation for compute nodes

Resource URI

Verb	URI
GET	/v1/diskmgmt/check_disks/?args={all,control,compute}

Example

JSON Request

```
GET /v1/diskmgmt/check_disks/?args=compute
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "add_as_spare_disks_results_list": [],
  "bad_disks_results_list": [],
  "fcfg_disks_results_list": [],
  "raid_results_list": [
    {
      "Num PDs": 4,
      "Num VDs": 1,
      "RAID health": "Opt",
      "RAID level": "RAID10",
      "RAID type": "HW",
      "VD health": "Opt1",
      "host": "i13-23",
      "role": "compute",
      "server": "15.0.0.6"
    },
    {
      "Num PDs": 4,
      "Num VDs": 1,
      "RAID health": "Opt",
      "RAID level": "RAID10",
      "RAID type": "HW",
      "VD health": "Opt1",
      "host": "i13-24",
      "role": "compute",
      "server": "15.0.0.10"
    }
  ],
  "rbld_disks_results_list": [],
  "spare_disks_results_list": []
}
```

Post a replace disk operation

Resource URI

Verb	URI
GET	/v1/diskmgmt/replace_disks/?args={all,control,compute}

Example

JSON Request

```
Get /v1/diskmgmt/replace_disks/?args=compute Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "add_as_spare_disks_results_list": [
    {
      "disk slot": "1",
      "host": "i13-21",
      "replace status": "Success",

```

```

        "role": "control",
        "server": "15.0.0.8"
    }
]
}

```

Get a check disk operation for a particular server

Resource URI

Verb	URI
GET	v1/diskmgmt/server/?server_list={server_list}&action=check_disks

Example

JSON Request

```

GET
/v1/diskmgmt/server/?server_list=i13-21,i13-23&action=check_disks
Accept: application/json

```

JSON Response

```

200 OK
Content-Type: application/json
{
  "add_as_spare_disks_results_list": [
    {
      "disk slot": "1",
      "disk state": " UGood",
      "host": "i13-21",
      "role": "control",
      "server": "15.0.0.8"
    }
  ],
  "bad_disks_results_list": [],
  "fcfg_disks_results_list": [],
  "raid_results_list": [
    {
      "Num PDs": 4,
      "Num VDs": 1,
      "RAID health": "NdAtn",
      "RAID level": "RAID10",
      "RAID type": "HW",
      "VD health": "Dgrd",
      "host": "i13-21",
      "role": "control",
      "server": "15.0.0.8"
    },
    {
      "Num PDs": 4,
      "Num VDs": 1,
      "RAID health": "Opt",
      "RAID level": "RAID10",
      "RAID type": "HW",
      "VD health": "Opt1",
      "host": "i13-23",
      "role": "compute",
      "server": "15.0.0.6"
    }
  ],
  "rbld_disks_results_list": [],
  "spare_disks_results_list": []
}

```

Perform a replace disk operation for a particular server

Resource URI

Verb	URI
GET	v1/diskmgmt/server/?server_list={server_list}&action={replace_disks}

Example

JSON Request

```
GET
/v1/diskmgmt/server?server_list=i13-21&action=replace_disks
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "add_as_spare_disks_results_list": [
    {
      "disk_slot": "1",
      "host": "i13-21",
      "replace_status": "Success",
      "role": "control",
      "server": "15.0.0.8"
    }
  ]
}
```

OSD Maintenance information

REST wrapper to query information about OSD on Pod storage nodes. This will return the OSD status information of all or a selection of OSDs available in the Pod.

End Point	Type	Valid Args	Valid Arg Values	Example
/osdmgmt	GET	None	None	/v1/osdmgmt/
/osdmgmt/check_osds	GET	None	Detail	/v1/osdmgmt/check_osds/
osdmgmt/server	GET	server_list, action	List of valid server names, check_osds, replace_osd, osd_name	/v1/osdmgmt/server/?server_list =svr1,svr2&action=check_osds /v1/osdmgmt/server/server_list=svr1&action =replace_osd&osd_name=osd_name

Get a OSD disk operation

Resource URI

Verb	URI
GET	/v1/osdmgmt

Example

JSON Request

```
GET
/v1/osdmgmt
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
{
  "bad_osds_results_list": [],
  "osd_details_results_list": [
    {
      "All OSD status": "All Good",
      "Num OSDs": 5,
      "OSD_detail": [
        {
          "OSD_id": 0,
          "OSD_journal": "/dev/sda4",
          "OSD_mount": "/var/lib/ceph/osd/ceph-0",
          "OSD_name": "osd.0",
          "OSD_path": "/dev/sdb1",
          "OSD_status": "up",
          "slot_id": 2
        },
        {
          "OSD_id": 3,
          "OSD_journal": "/dev/sda5",
          "OSD_mount": "/var/lib/ceph/osd/ceph-3",
          "OSD_name": "osd.3",
          "OSD_path": "/dev/sdc1",
          "OSD_status": "up",
          "slot_id": 3
        },
        {
          "OSD_id": 6,
          "OSD_journal": "/dev/sda6",
          "OSD_mount": "/var/lib/ceph/osd/ceph-6",
          "OSD_name": "osd.6",
          "OSD_path": "/dev/sdd1",
          "OSD_status": "up",
          "slot_id": 4
        },
        {
          "OSD_id": 9,
          "OSD_journal": "/dev/sda7",
          "OSD_mount": "/var/lib/ceph/osd/ceph-9",
          "OSD_name": "osd.9",
          "OSD_path": "/dev/sde1",
          "OSD_status": "up",
          "slot_id": 5
        },
        {
          "OSD_id": 12,
          "OSD_journal": "/dev/sda8",
          "OSD_mount": "/var/lib/ceph/osd/ceph-12",
          "OSD_name": "osd.12",
          "OSD_path": "/dev/sdf1",
          "OSD_status": "up",
          "slot_id": 6
        }
      ]
    },
    {
      "host": "i13-27-test",
      "role": "block_storage",
    }
  ]
}
```

```

    "server": "15.0.0.4"
  },
  {
    "All OSD status": "All Good",
    "Num OSDs": 5,
    "OSD_detail": [
      {
        "OSD_id": 1,
        "OSD_journal": "/dev/sda4",
        "OSD_mount": "/var/lib/ceph/osd/ceph-1",
        "OSD_name": "osd.1",
        "OSD_path": "/dev/sdb1",
        "OSD_status": "up",
        "slot_id": 2
      },
      {
        "OSD_id": 4,
        "OSD_journal": "/dev/sda5",
        "OSD_mount": "/var/lib/ceph/osd/ceph-4",
        "OSD_name": "osd.4",
        "OSD_path": "/dev/sdc1",
        "OSD_status": "up",
        "slot_id": 3
      },
      {
        "OSD_id": 7,
        "OSD_journal": "/dev/sda6",
        "OSD_mount": "/var/lib/ceph/osd/ceph-7",
        "OSD_name": "osd.7",
        "OSD_path": "/dev/sdd1",
        "OSD_status": "up",
        "slot_id": 4
      },
      {
        "OSD_id": 10,
        "OSD_journal": "/dev/sda7",
        "OSD_mount": "/var/lib/ceph/osd/ceph-10",
        "OSD_name": "osd.10",
        "OSD_path": "/dev/sde1",
        "OSD_status": "up",
        "slot_id": 5
      },
      {
        "OSD_id": 13,
        "OSD_journal": "/dev/sda8",
        "OSD_mount": "/var/lib/ceph/osd/ceph-13",
        "OSD_name": "osd.13",
        "OSD_path": "/dev/sdf1",
        "OSD_status": "up",
        "slot_id": 6
      }
    ],
    "host": "i13-25",
    "role": "block_storage",
    "server": "15.0.0.11"
  },
  {
    "All OSD status": "All Good",
    "Num OSDs": 5,
    "OSD_detail": [
      {
        "OSD_id": 2,
        "OSD_journal": "/dev/sda4",
        "OSD_mount": "/var/lib/ceph/osd/ceph-2",

```

```

        "OSD_name": "osd.2",
        "OSD_path": "/dev/sdb1",
        "OSD_status": "up",
        "slot_id": 2
    },
    {
        "OSD_id": 5,
        "OSD_journal": "/dev/sda5",
        "OSD_mount": "/var/lib/ceph/osd/ceph-5",
        "OSD_name": "osd.5",
        "OSD_path": "/dev/sdc1",
        "OSD_status": "up",
        "slot_id": 3
    },
    {
        "OSD_id": 8,
        "OSD_journal": "/dev/sda6",
        "OSD_mount": "/var/lib/ceph/osd/ceph-8",
        "OSD_name": "osd.8",
        "OSD_path": "/dev/sdd1",
        "OSD_status": "up",
        "slot_id": 4
    },
    {
        "OSD_id": 11,
        "OSD_journal": "/dev/sda7",
        "OSD_mount": "/var/lib/ceph/osd/ceph-11",
        "OSD_name": "osd.11",
        "OSD_path": "/dev/sde1",
        "OSD_status": "up",
        "slot_id": 5
    },
    {
        "OSD_id": 14,
        "OSD_journal": "/dev/sda8",
        "OSD_mount": "/var/lib/ceph/osd/ceph-14",
        "OSD_name": "osd.14",
        "OSD_path": "/dev/sdf1",
        "OSD_status": "up",
        "slot_id": 6
    }
],
"host": "i13-26",
"role": "block_storage",
"server": "15.0.0.9"
}
]
}

```

Perform a check OSD operation for a particular server

Resource URI

Verb	URI
GET	/v1/osdmgmt/server/?server_list={server_list}&action={check_osds}

Example

JSON Request

```

GET
/v1/diskmgmt/server/?server_list=i13-26&action=check_osds
Accept: application/json

```

JSON Response

200 OK

Content-Type: application/json

```

{
  "bad_osds_results_list": [],
  "osd_details_results_list": [
    {
      "All OSD status": "All Good",
      "Num OSDs": 5,
      "OSD_detail": [
        {
          "OSD_id": 2,
          "OSD_journal": "/dev/sda4",
          "OSD_mount": "/var/lib/ceph/osd/ceph-2",
          "OSD_name": "osd.2",
          "OSD_path": "/dev/sdb1",
          "OSD_status": "up",
          "slot_id": 2
        },
        {
          "OSD_id": 5,
          "OSD_journal": "/dev/sda5",
          "OSD_mount": "/var/lib/ceph/osd/ceph-5",
          "OSD_name": "osd.5",
          "OSD_path": "/dev/sdc1",
          "OSD_status": "up",
          "slot_id": 3
        },
        {
          "OSD_id": 8,
          "OSD_journal": "/dev/sda6",
          "OSD_mount": "/var/lib/ceph/osd/ceph-8",
          "OSD_name": "osd.8",
          "OSD_path": "/dev/sdd1",
          "OSD_status": "up",
          "slot_id": 4
        },
        {
          "OSD_id": 11,
          "OSD_journal": "/dev/sda7",
          "OSD_mount": "/var/lib/ceph/osd/ceph-11",
          "OSD_name": "osd.11",
          "OSD_path": "/dev/sde1",
          "OSD_status": "up",
          "slot_id": 5
        },
        {
          "OSD_id": 14,
          "OSD_journal": "/dev/sda8",
          "OSD_mount": "/var/lib/ceph/osd/ceph-14",
          "OSD_name": "osd.14",
          "OSD_path": "/dev/sdf1",
          "OSD_status": "up",
          "slot_id": 6
        }
      ]
    },
    {
      "host": "i13-26",
      "role": "block_storage",
      "server": "15.0.0.9"
    }
  ]
}

```

Perform a replace OSD operation for a particular server

Resource URI

Verb	URI
GET	/v1/osdmgmt/server/?server_list={server_name}&action=replace_osd&osd_name={osd_name}

Example

JSON Request

```
GET
/v1/diskmgmt/server/?server_list=i13-25&action=replace_osd&osd_name=osd.10
Accept: application/json
```

JSON Response

```
200 OK
Content-Type: application/json
```

```
{
  'osd_replace_details_results_list': [
    {
      'hdd_slot': 5,
      'host_name': 'i13-25',
      'journal_mnt': '/dev/sda4',
      'new_dev_uuid': 'UUID=94480b3e-5698-4d6a-b715-0613be41cff5',
      'new_mount': '/var/lib/ceph/osd/ceph-10',
      'new_osd_id': 10,
      'new_path': '/dev/sde1',
      'old_osd_id': 10,
      'status_msg': 'Successfully deleted, removed and replaced OSD
osd.10 from server i13-25'
    }
  ]
}
```



CHAPTER 3

Monitoring Cisco NFVI Performance

The following topics tell you how to display logs to monitor Cisco VIM performance.

- [Logging and Monitoring in Cisco NFVI, on page 83](#)
- [Displaying Cisco VIM Log Files Using the CLI, on page 85](#)
- [Logging Into the Kibana Dashboard, on page 88](#)
- [Rotation of the Cisco VIM Logs, on page 97](#)
- [Network Performance Test with NFVBench, on page 97](#)

Logging and Monitoring in Cisco NFVI

Cisco VIM uses a combination of open source tools to collect and monitor the Cisco OpenStack services including Elasticsearch, Fluentd, and the Kibana dashboard (EFK).

In VIM, we have moved our platform to use Fluentd, instead of logstash. However, to maintain backwards compatibility, the code, and documentation refers to ELK, instead of EFK at various places. In VIM, these two acronyms are interchangeable, however it refers to the presence of EFK in the offering. OpenStack services that followed by EFK include:

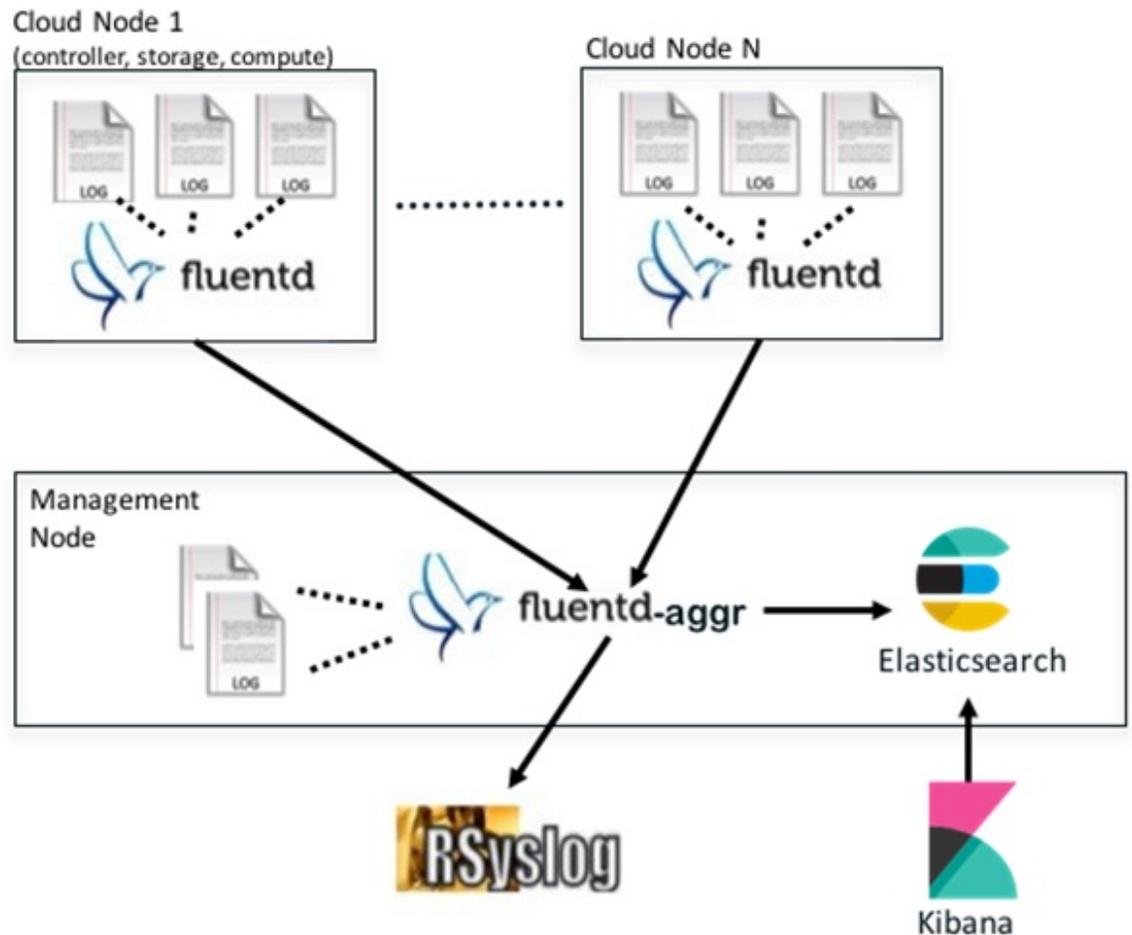
- **MariaDB**—A relational database management system which is based on MySQL. All the OpenStack components store their data in MariaDB.
- **HAProxy**—HAProxy is a free open source software that provides a high-availability load balancer, and proxy server for TCP and HTTP-based applications that spreads requests across multiple servers.
- **Keystone**—Keystone is an OpenStack project that provides identity, token, catalog, and policy services for use specifically by projects in the OpenStack.
- **Glance**—An OpenStack project that allows you to upload and discover data assets that are meant for use with other services.
- **Neutron**—An OpenStack project that provides the network connectivity between interface devices, such as vNICs, managed by other OpenStack services, such as Nova.
- **Nova**—An OpenStack project that is designed to provide massively scalable, on demand, self-service access to compute resources.
- **HTTP**—The Apache HTTP server Project, an effort to develop and maintain an open-source HTTP server.

- Cinder—An OpenStack block storage service that is designed to present storage resources to the users that are consumed by the OpenStack compute project (Nova).
- Memcached—A general purpose distributed memory caching system.
- CloudPulse—Is an OpenStack tool that checks the health of the cloud. CloudPulse includes operator and end-point tests.
- Heat—The main OpenStack Orchestration program. Heat implements an orchestration engine to launch multiple composite cloud applications that is based on text file templates.
- Other OpenStack services—RabbitMQ, Ceph, Open vSwitch, Linux bridge, Neutron VTS (optional), and others.
- VMTP—Integrated control and data plane log for testing the cloud.
- NFVBench—Network performance benchmarking tool.

A Fluentd container resides on each control, compute, and storage nodes. They forward log to the Fluentd-aggr server residing on the management node.

The following figure shows a high-level schematic of the Fluent service assurance architecture.

Figure 8: EFK Service Assurance Architecture



The EFK flow includes:

- Fluentd extracts the relevant data from the logs and tags them so that Kibana can use it later to display useful information about those logs.
- Fluentd sends the logs from all the compute, controller, and storage nodes to the Fluentd-aggr server on the management node.
- Fluentd-aggr in the management node sends the structured logs into the Elasticsearch database.
- Elasticsearch stores the data, indexes it, and supports fast queries against a large amount of log data.
- Kibana visualizes the data that is stored in Elasticsearch using a custom dashboard. You can also add filters to the data to visualize interesting fragments of the log data.

Displaying Cisco VIM Log Files Using the CLI

Cisco VIM log file location depends on the node and log type. Installer logs are found in the management node under the `/var/log/mercury/<install_uid>/` directory. The last 20 log directories are tarred and kept in this directory. These files contain logs related to bootstrap, build orchestration, baremetal, common setup, and OpenStack orchestration.

If the installer fails, look at the last tar.gz file for logs, for example:

```
[root@mgmtnode mercury]# ls -lrt
total 20
drwxr-xr-x. 2 root root 80 Jul 19 23:42 573f2b7f-4463-4bfa-b57f-98a4a769aced
drwxr-xr-x. 2 root root 4096 Jul 20 03:29 installer
drwxr-xr-x. 2 root root 79 Jul 20 03:29 e9117bc5-544c-4bda-98d5-65bffa56a18f
drwxr-xr-x. 2 root root 79 Jul 20 04:54 36cdf8b5-7a35-4e7e-bb79-0cfb1987f550
drwxr-xr-x. 2 root root 79 Jul 20 04:55 bd739014-fdf1-494e-adc0-98b1fba510bc
drwxr-xr-x. 2 root root 79 Jul 20 04:55 e91c4a6c-ae92-4fef-8f7c-cafa9f5dc1a3
drwxr-xr-x. 2 root root 79 Jul 20 04:58 1962b2ba-ff15-47a6-b292-25b7fb84cd28
drwxr-xr-x. 2 root root 79 Jul 20 04:59 d881d453-f6a0-448e-8873-a7c51d8cc442
drwxr-xr-x. 2 root root 78 Jul 20 05:04 187a15b6-d425-46a8-a4a2-e78b65e008b6
drwxr-xr-x. 2 root root 4096 Jul 20 06:47 d0346cdd-5af6-4058-be86-1330f7ae09d1
drwxr-xr-x. 2 root root 79 Jul 20 17:09 f85c8c6c-32c9-44a8-b649-b63fdb11a79a
drwxr-xr-x. 2 root root 67 Jul 20 18:09 179ed182-17e4-4f1f-a44d-a3b6c16cf323
drwxr-xr-x. 2 root root 68 Jul 20 18:13 426cb05f-b1ee-43ce-862d-5bb4049cc957
drwxr-xr-x. 2 root root 68 Jul 20 18:13 1d2eec9d-f4d8-4325-9eb1-7d96d23e30fc
drwxr-xr-x. 2 root root 68 Jul 20 18:13 02f62a2f-3f59-46a7-9f5f-1656b8721512
drwxr-xr-x. 2 root root 68 Jul 20 18:14 c7417be9-473e-49da-b6d0-d1ab8fb4b1fc
drwxr-xr-x. 2 root root 68 Jul 20 18:17 b4d2077b-c7a9-46e7-9d39-d1281fba9baf
drwxr-xr-x. 2 root root 68 Jul 20 18:35 21972890-3d45-4642-b41d-c5fadfeba21a
drwxr-xr-x. 2 root root 80 Jul 20 19:17 d8b1b54c-7fc1-4ea6-83a5-0e56ff3b67a8
drwxr-xr-x. 2 root root 80 Jul 20 19:17 23a3cc35-4392-40bf-91e6-65c62d973753
drwxr-xr-x. 2 root root 80 Jul 20 19:17 7e831ef9-c932-4b89-8c81-33a45ad82b89
drwxr-xr-x. 2 root root 80 Jul 20 19:18 49ea0917-f9f4-4f5d-82d9-b86570a02dad
drwxr-xr-x. 2 root root 80 Jul 20 19:18 21589a61-5893-4e30-a70e-55ad0dc2e93f
drwxr-xr-x. 2 root root 80 Jul 20 19:22 6ae6d136-7f87-4fc8-92b8-64cd542495bf
drwxr-xr-x. 2 root root 4096 Jul 20 19:46 1c6f4547-c57d-4dcc-a405-ec509306ee25
drwxr-xr-x. 2 root root 68 Jul 20 21:20 c6dcc98d-b45b-4904-a217-d25001275c85
drwxr-xr-x. 2 root root 68 Jul 20 21:40 ee58d5d6-8b61-4431-9f7f-8cab2c331637
drwxr-xr-x. 2 root root 4096 Jul 20 22:06 243cb0f8-5169-430d-a5d8-48008a00d5c7
drwxr-xr-x. 2 root root 4096 Jul 20 22:16 188d53da-f129-46d9-87b7-c876b1aea70c
```

Cisco VIM autobackup logs are found in the following location:

```
# CVIM autobackup logs (auto-backup enabled by default)
/var/log/mercury/autobackup_2.2.x_2018-03-19_15-11-10.log

# Cobbler apache log (may be needed for PXE troubleshooting)
/var/log/cobblerhttpd/access_log
/var/log/cobblerhttpd/error_log

# VMTP logs
/var/log/vmtp/vmtp.log
```

Cisco VIM RestAPI log location

```
# CVIM RestAPI logs
/var/log/mercury_restapi/restapi.log

# CIM RestAPI apache logs (TCP port 8445)
/var/log/httpd/mercury_access.log
/var/log/httpd/mercury_error.log

# CIM RestAPI log-directory logs (TCP port 8008)
/var/log/httpd/access_log
/var/log/httpd/error_log
```

EFK log location

```
# Elasticsearch-fluentd-Kibana
/var/log/elasticsearch/
/var/log/fluentd-aggr/
/var/log/kibana/
/var/log/curator/
```

```
# HAProxy TLS certificate expiration check
/var/log/curator/certchecker.log
```

Viewing Cisco VIM Logs

```
# list logs sorted reverse on time
ls -lrt /var/log/mercury/
# untar logs
tar xvzf /var/log/mercury/<UUID>/mercury_install_2018-3-20_10-2.tar.gz -C /tmp/
```

Cisco VIM Configuration Files

```
# example configuration files
/root/openstack-configs/setup_data.yaml.B_Series_EXAMPLE
/root/openstack-configs/setup_data.yaml.C_Series_EXAMPLE

# system maintained setup files - do not modify directly
# always supply user copy of setup_data.yaml
# when using ciscovim client
/root/openstack-configs/setup_data.yaml

# system inventory in pretty format
/root/openstack-configs/mercury_servers_info

# passwords store
/root/openstack-configs/secrets.yaml

# openstack configuration file
/root/openstack-configs/openstack_config.yaml

# RestAPI password
/opt/cisco/ui_config.json

# Insight password
/opt/cisco/insight/secrets.yaml
```

Enabling debug logs for certain OpenStack Services

```
# openstack config file
/root/openstack-configs/openstack_config.yaml

# help
ciscovim help

# list openstack keys
ciscovim list-openstack-configs

# help on reconfigure sub-command
ciscovim help reconfigure

# how to execute subcommand, example below
# important note: reconfigure requires a maintenance window
ciscovim reconfigure --setopenstackconfig KEYSTONE_DEBUG_LOGGING,CINDER_DEBUG_LOGGING
```

On controller and compute nodes, all services are run within their respective Docker™ containers.

To list the Docker containers in the node, execute the following:

```
[root@control-server-2 ~]# docker ps -a
CONTAINER ID          IMAGE                                     PORTS                NAMES                COMMAND
258b2cald46a         172.31.228.164:5000/mercury-rhel7-osp8/nova-scheduler:4780
"/usr/bin/my_init /no" 25 minutes ago     Up 25 minutes       novascheduler_4780
ffe70809bbe0         172.31.228.164:5000/mercury-rhel7-osp8/nova-novncproxy:4780
"/usr/bin/my_init /st" 25 minutes ago     Up 25 minutes       novanovncproxy_4780
```

```
12b92bcb9dc0          172.31.228.164:5000/mercury-rhel7-osp8/nova-consoleauth:4780
"/usr/bin/my_init /st" 26 minutes ago   Up 26 minutes
```

.....

```
novaconsoleauth_4780
7295596f5167          172.31.228.164:5000/mercury-rhel7-osp8/nova-api:4780
"/usr/bin/my_init /no" 27 minutes ago   Up 27 minutes          novaapi_4780
```

To view the Docker logs of any container, execute the following on the corresponding host:

```
ls -l /var/log/<service_name>/<log_filename>
e.g. ls -l /var/log/keystone/keystone.log
```

To get into a specific container, execute the following commands:

```
[root@control-server-2 ~]# alias | grep container
      root@control-server-2 ~]# source /root/.bashrc
#execute the alias:
      [root@control-server-2 ~]# novaapi
novaapi_4761 [nova@control-server-2 /]$
novaapi_4761 [nova@control-server-2 /]$ exit
exit
```

If the Docker status indicates a container is down (based on output of “docker ps -a”), collect the Docker service logs as well:

```
cd /etc/systemd/system/multi-user.target.wants/
ls docker* # get the corresponding service name from the output
systemctl status <service_name> -n 1000 > /root/filename # redirects the output to the file
```

For storage nodes running Ceph, execute the following to check the cluster status:

```
ceph -v # on monitor nodes (controller), show's ceph version
ceph -s # on monitor nodes (controller), show cluster status
ceph osd lspools #on monitor nodes (controller),list pools
ceph mon stat # summarize monitor status
ceph-disk list # on OSD / storage nodes; List disks, partitions, and Ceph OSDs
rbd list images # on monitor nodes (controller); dump list of image snapshots
rbd list volumes # on monitor nodes (controller); dump list of volumes
```

Logging Into the Kibana Dashboard

Kibana is an open source data visualization platform that you can use to explore Cisco VIM logs.

To log into the Kibana dashboard:

Step 1 Using a terminal client, use SSH to log into your management node and enter the password to login.

The following command shows the management node has an IP address of 17.0.0.2:

```
# ssh root@17.0.0.2
root@17.0.0.2's password
```

Step 2 In the SSH terminal session, locate the line containing ELK_PASSWORD in /root/installer-{tag id}/openstack-configs/secrets.yaml. Note the value of the ELK_PASSWORD. It is used in Step 4.

```
cat /root/installer-{tag-id}/openstack-configs/secrets.yaml
...
ELK_PASSWORD: <note this value>
...
```

Step 3 Using your web browser, navigate to `http://<management_node_ip_address>:5601`.

Step 4 When prompted, log in with the following credentials:

User Name: admin

Password: <value of ELK_PASSWORD from Step 2>

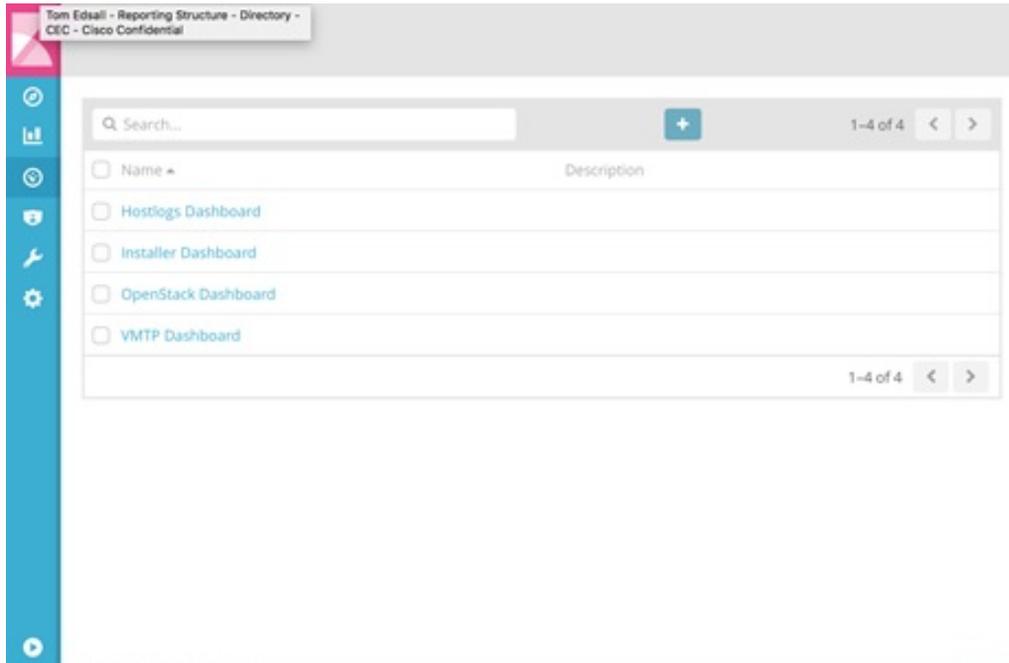
The Kibana dashboard appears allowing you to display the Cisco VIM service and installer logs.

Figure 9: Editing New Dashboard



Step 5 Navigate to the dashboards by clicking **Dashboard** menu bar and choose the desired dashboard. It is not recommended to use visualize/Timelion/DevTools or Management options on the left side.

Figure 10: Lists of Dashboards

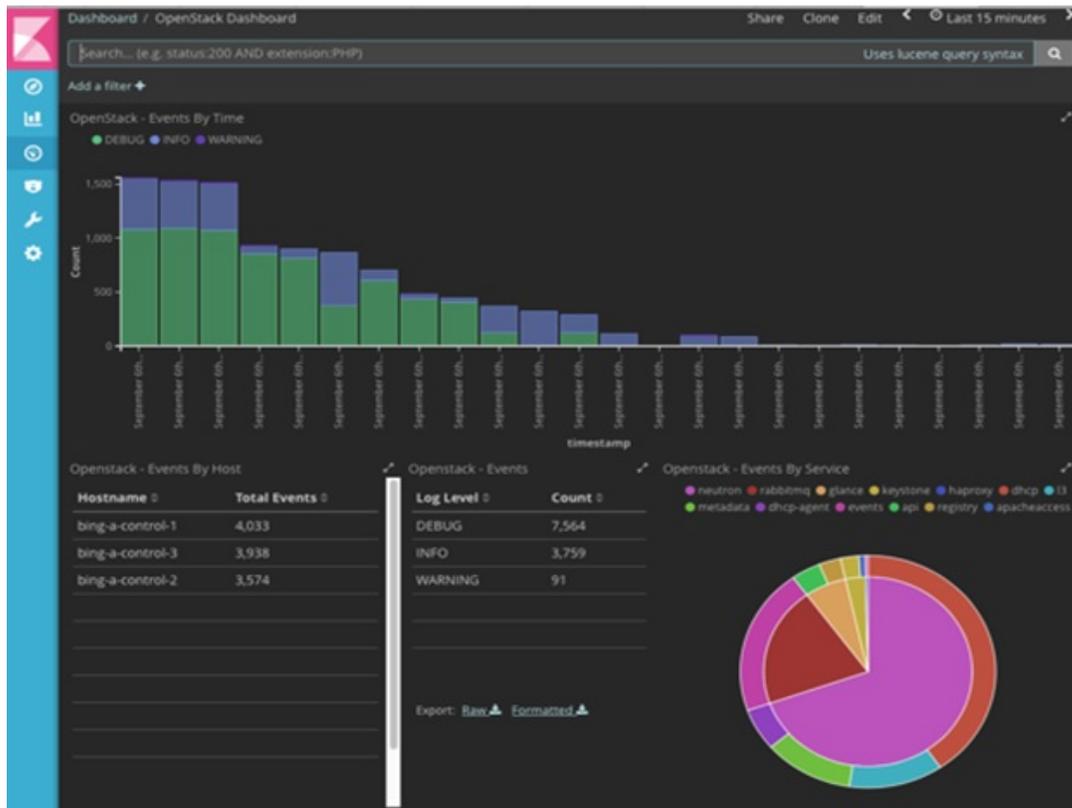


The following are the list of dashboards supported:

- **Hostlogs Dashboard:** Provides log information of the system for the cloud nodes. This displays entries from the host logs-* index in Elasticsearch. It contains the log from /var/log/messages file on each server.
- **Installer Dashboard:** Provides information about the management node and the installation process. It can only read uncompressed files. Hence, it reads the files prior to the cloud installation. This displays entries from the installer-* index in Elasticsearch.
- **OpenStack Dashboard:** (openstack-* index) Provides log information about all the OpenStack processes. This displays entries from the openstack-* index in Elasticsearch.
- **VMTP Dashboard:** Provides log information about the VMTP runs performed against the cloud. It displays entries from the vmtp-* index in Elasticsearch.

For Example: if you click on the OpenStack dashboard link the following screen appears.

Figure 11: OpenStack Dashboard



You can switch on from one dashboard to another by selecting the appropriate dashboard from the right top bar menu.

All dashboards have generic and specific fields.

The generic ones are:

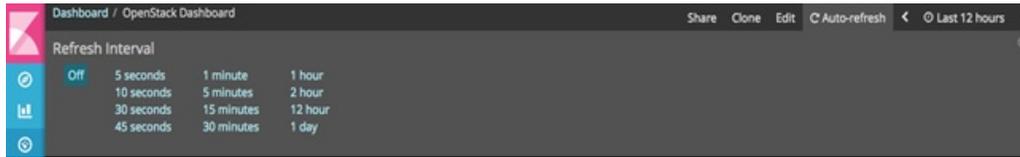
- Title: Title is seen at the top left of the page. Title shows which dashboard is being displayed. For Example: OpenStack Dashboard.
- Time: Time is seen at the top right of the page. Time indicates the time schedule for the log information. You can modify the time to indicate absolute, relative time in the past or specify automatically refresh rates.
- Search bar: Search bar is an input field where you can enter a query in the Lucene syntax format to filter the logs by specific fields (which depend on the fields for the index being selected)
- Add a filter tab: Use this tab to introduce filters graphically.

For more information on using Kibana, see the *Kibana documentation* (Version 5.5.1).

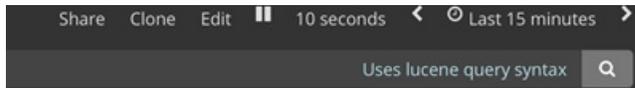
Cisco VIM stores the OpenStack logs in Elasticsearch. The Elasticsearch snapshots all the indices (where the data is stored) which are rotated on a periodic basis. You may not see the older data in Kibana if the data is rotated out and/or deleted.

Logs keep being visualized in Kibana as they are being updated in Elasticsearch on the Discover tab. To debug something on kibana, you can program the Kibana dashboard to auto-refresh at specific intervals (by default is off). To enable auto-refresh, click the date at the top right corner of the dashboard and click Auto-refresh to configure the desired value.

Figure 12: Auto-Refresh



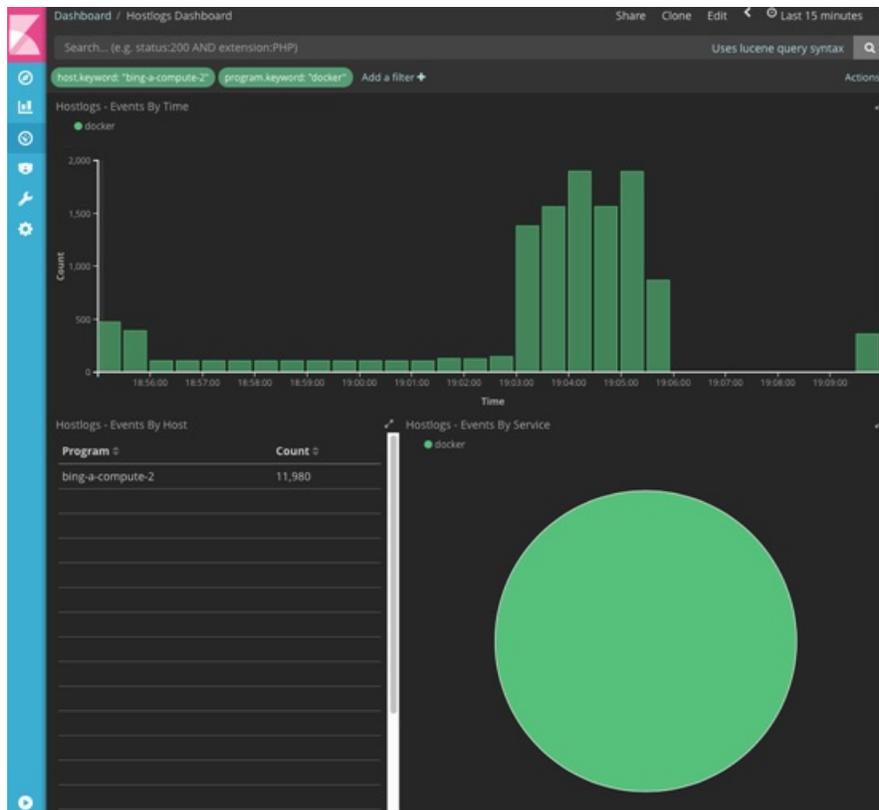
User can click the play/pause button on the top navigator bar to continue/pause the refreshing of logs events:



a) Few examples on usage of filters in Openstack dashboard to gather useful information

- On the Hostlogs Dashboard, in the Events by Host panel, choose a hostname and click the + or - symbol that appears close to the hostname to include or exclude that server from the filter. Then, click the desired slice on the Events By Service panel to add the docker service to the section.
- Under the Search field, you can see the included sections in green and excluded sections in red.

Figure 13: Hostlogs Dashboard



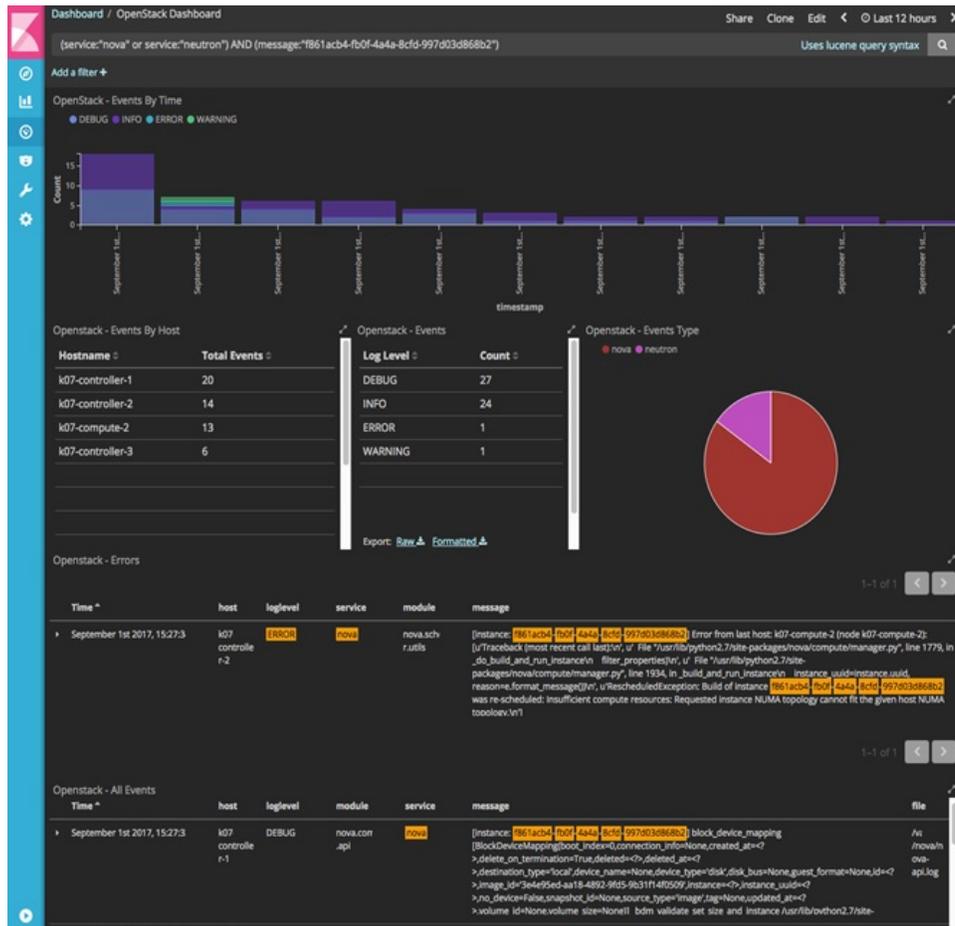
Time	logdate	host	program	message
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.342 8 ERROR oslo.messaging._drivers.impl_rabbit [] Failed to consume message from queue: (0, 0) (403) ACCESS_REFUSED - Login was refused using authentication mechanism AMQPPLAIN. For details see the broker logfile.
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.467 8 DEBUG neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent [req-809cb46-3eaa-492f-a352-481b34f82210 - ... -] Agent rpc_loop - iteration:1397 started rpc_loop /usr/lib/python2.7/site-packages/neutron/plugins/ml2/drivers/openvswitch/agent/ovs_neutron_agent.py:1965
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.472 8 DEBUG neutron.agent.linux.utils [req-809cb46-3eaa-492f-a35181b34f82210 - ... -] Running command: [ps, -ppid, '85', '0', 'pid='] create_process /usr/lib/python2.7/site-packages/neutron/agent/linux/utils.py:89
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.095 8 ERROR oslo.messaging._drivers.impl_rabbit [] [5d14257f-d6f1-4b33-8532-4d00421b66ea] AMQP server on 10.23.222.122:5672 is unreachable: <AMQPError: unknown error>. Trying again in 1 seconds. Client port: None
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.096 8 ERROR oslo.messaging._drivers.impl_rabbit [] [9521d890-240f-46a4-9f9b-5037ad5206b6] AMQP server on 10.23.222.122:5672 is unreachable: <AMQPError: unknown error>. Trying again in 1 seconds. Client port: None
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.574 8 DEBUG neutron.agent.linux.utils [req-809cb46-3eaa-492f-a35181b34f82210 - ... -] Exit code: 0 execute /usr/lib/python2.7/site-packages/neutron/agent/linux/utils.py:150
September 6th 2017, 19:09:4	-	bing-compute-2	docker	2017-09-07 02:09:44.470 8 DEBUG neutron.plugins.ml2.drivers.openvswitch.agent.openflow.native.ofswitch [req-809cb46-3eaa-492f-a352-481b34f82210 - ... -] ofctl request version=0x4,mgmt_type=0x12,mgmt_seq=0x38,vid=0x496923cb,DPFFlowStatsRequest(cookie=0,cookie_mask=0,flags=0,match=OFFMatchOam,fields=(),out_group=4294967295,out_port=4294967295,table_id=23,type=1) result IDFFFlowStatsReolvbody=

b) To know the log events in the Openstack for a given VM by writing the filter directly on the Search field:

Note The uuid of the VM can be identified by executing `openstack nova list` or looking at the horizon website.

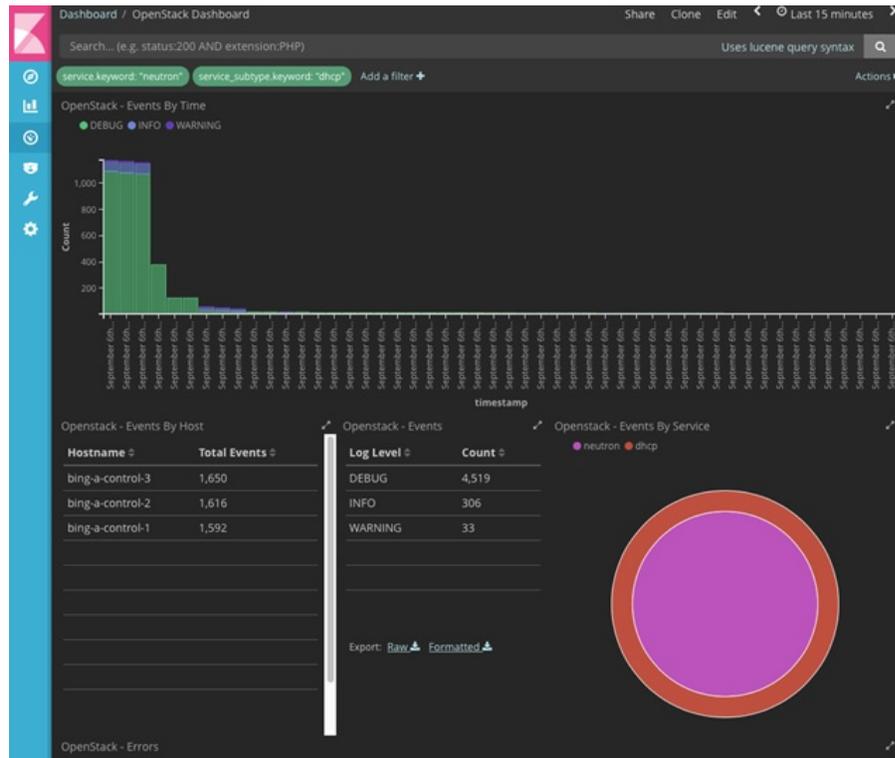
- In the Search field which is on top of the Dashboard, write the following Lucene query: `(service: nova and service: neutron) AND (message:<uuid>)` here, `<uuid>` is the number got from Horizon or nova list for the identifier of the instance VM.

Figure 14: Search Query Page

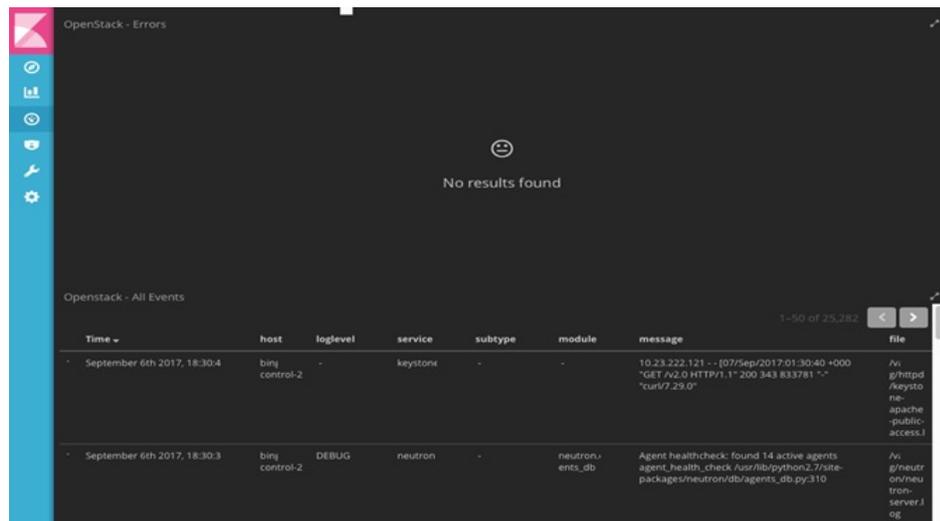


- For example, if the user wants to know the DHCP events of the Openstack Neutron add filters by clicking outer circle of pie chart:
 - On the OpenStack Dashboard, the Openstack - Events By Service panel has a pie chart with the inner section for the services and the outer sections for the service_subtypes. To add filters for selecting all the events in a service (for example, neutron), click on the inner section of the pie. To add filters for selecting the service_subtypes (for example, dhcp), click on the outer circle of the pie.

Figure 15: Events by Service



- Note: You can scroll down the OpenStack Dashboard to see the OpenStack - Errors and the OpenStack - Events panel. The OpenStack - Errors panel displays the error messages. If there are no errors, the **No results found** message is displayed.

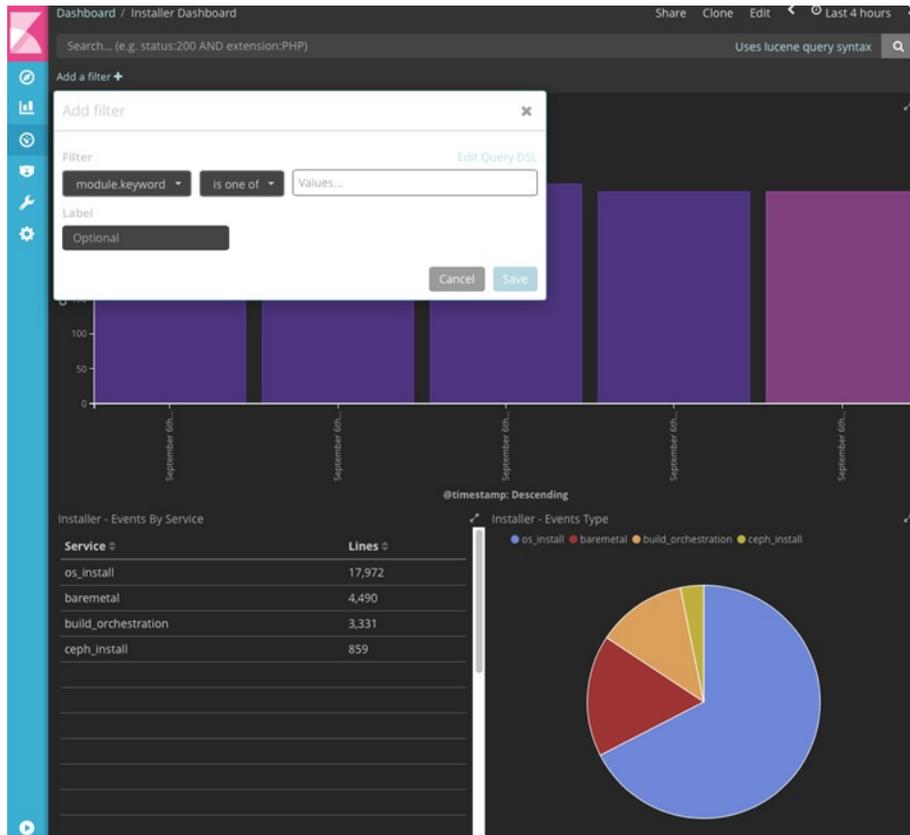


- Without knowing the Lucene Syntax, you can set the filter criteria in the **Search** field using the **Add a filter +** option.

Following are the steps to add a filter:

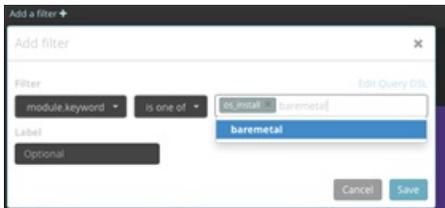
- Click Add a filter (+).
- Set the filter criteria by choosing appropriate label and operators from the drop-down lists, and entering keywords and click Save.

Figure 16: Add Filters Page



Set the filter criteria by choosing appropriate label and operators from the drop-down lists, and entering keywords.

Figure 17: Choosing Appropriate Labels



Rotation of the Cisco VIM Logs

Cisco VIM stores all logs in Elasticsearch. Elasticsearch indices are rotated on a periodic basis to prevent the disk space overflow by creating snapshots. The following lists show the Snapshots that are defined in `openstack_config.yaml`:

```
# vi ~/openstack-configs/openstack_config.yaml
...
# Elk rotation parameters
elk_rotation_frequency: "monthly" # Available: "daily", "weekly", "fortnightly", "monthly"
elk_rotation_size: 2 # Unit is in Gigabytes (float is allowed)
elk_rotation_del_older: 10 # Delete older than 10 units (where units depends on the
value set on elk_rotation_frequency)
...
```

You can change the frequency of the rotation by changing the values. For more information on how to set the Elasticsearch parameters through VIM API or CLI, refer to the section *Reconfiguring Passwords and OpenStack Configurations*.

Cisco VIM uses the open source Elasticsearch Curator tool to manage the Elasticsearch indices and snapshots. For more information about Elasticsearch handles snapshots, look at the official information on Elastic.co (Version 5.4) <https://www.elastic.co/guide/en/elasticsearch/client/curator/5.4/index.html>.

Network Performance Test with NFVBench

NFVBench is a network performance benchmarking tool integrated with Cisco VIM. For more details, refer to NFVBench section of *Chapter 1* in the admin guide for details.



CHAPTER 4

Managing Cisco NFVI Security

The following topics describe Cisco NFVI network and application security and best practices.

- [Verifying Management Node Network Permissions, on page 99](#)
- [Verifying Management Node File Permissions, on page 100](#)
- [Viewing Administrator Access Attempts, on page 100](#)
- [Verifying SELinux, on page 101](#)
- [Validating Port Listening Services, on page 101](#)
- [Validating Non-Root Users for OpenStack Services, on page 102](#)
- [Verifying Password Strength, on page 102](#)
- [Reconfiguring Passwords and OpenStack Configurations, on page 103](#)
- [Enabling NFVIMON Post Pod Install, on page 106](#)
- [Reconfiguring CIMC Password on an Existing Install, on page 108](#)
- [Increasing Provider and Tenant VLAN Ranges, on page 109](#)
- [Fernet Key Operations, on page 109](#)
- [Managing Certificates, on page 110](#)
- [Reconfiguring TLS Certificates, on page 110](#)
- [Enabling Keystone v3 on an Existing Install, on page 111](#)

Verifying Management Node Network Permissions

The Cisco NFVI management node stores sensitive information related to Cisco NFVI operations. Access to the management node can be restricted to requests coming from IP addresses known to be used by administrators. The administrator source networks is configured in the setup file, under **[NETWORKING]** using the **admin_source_networks** parameter.

To verify this host based firewall setting, log into the management node as an admin user and list the rules currently enforces by iptables. Verify that the source networks match the values configured. If no source networks have been configured, then all source traffic is allowed. However, note that only traffic destined to ports with known admin services is allowed to pass. The **admin_source_networks** value can be set at install time or changed through a reconfigure.

```
[root@control-server-1 ~]# iptables -list
Chain INPUT (policy ACCEPT)
target     prot opt source                destination
ACCEPT    icmp -- anywhere              anywhere
ACCEPT    tcp  -- 10.0.0.0/8            anywhere          tcp dpt:ssh
ACCEPT    tcp  -- 172.16.0.0/12         anywhere          tcp dpt:ssh
```

```

ACCEPT      tcp  --  10.0.0.0/8          anywhere          tcp dpt:https
ACCEPT      tcp  --  172.16.0.0/12      anywhere          tcp dpt:https
ACCEPT      tcp  --  10.0.0.0/8          anywhere          tcp dpt:4979
ACCEPT      tcp  --  172.16.0.0/12      anywhere          tcp dpt:4979
ACCEPT      tcp  --  10.0.0.0/8          anywhere          tcp dpt:esmagent
ACCEPT      tcp  --  172.16.0.0/12      anywhere          tcp dpt:esmagent
ACCEPT      tcp  --  10.0.0.0/8          anywhere          tcp dpt:8008
ACCEPT      tcp  --  172.16.0.0/12      anywhere          tcp dpt:8008
ACCEPT      tcp  --  10.0.0.0/8          anywhere          tcp dpt:copy
ACCEPT      tcp  --  172.16.0.0/12      anywhere          tcp dpt:copy
ACCEPT      tcp  --  10.0.0.0/8          anywhere          tcp dpt:22250
ACCEPT      tcp  --  172.16.0.0/12      anywhere          tcp dpt:22250
ACCEPT      all  --  anywhere           anywhere          state RELATED,ESTABLISHED
DROP        all  --  anywhere           anywhere

```

Verifying Management Node File Permissions

The Cisco NFVI management node stores sensitive information related to Cisco NFVI operations. These files are secured by strict file permissions. Sensitive files include `secrets.yaml`, `openrc`, `*.key`, and `*.pem`. To verify the file permissions, log into the management node as an admin user and list all of the files in the `~/openstack-configs/` directory. Verify that only the owner has read and write access to these files. For example:

```

[root@control-server-1 ~]# ls -l ~/openstack-configs
total 172
-rw-----. 1 root root 3272 Jun 21 17:57 haproxy.key
-rw-----. 1 root root 5167 Jun 21 17:57 haproxy.pem
-rw-----. 1 root root 223 Aug 8 18:09 openrc
-rw-----. 1 root root 942 Jul 6 19:44 secrets.yaml

[...]

```

Viewing Administrator Access Attempts

As the UCS servers are part of the critical Cisco NFVI infrastructure, Cisco recommends monitoring administrator login access periodically.

To view the access attempts, use the `journalctl` command to view the log created by `ssh`. For example:

```

[root@control-server-1 ~]# journalctl -u sshd
-- Logs begin at Tue 2016-06-21 17:39:35 UTC, end at Mon 2016-08-08 17:25:06 UTC. --
Jun 21 17:40:03 hh23-12 systemd[1]: Started OpenSSH server daemon.
Jun 21 17:40:03 hh23-12 systemd[1]: Starting OpenSSH server daemon...
Jun 21 17:40:03 hh23-12 sshd[2393]: Server listening on 0.0.0.0 port 22.
Jun 21 17:40:03 hh23-12 sshd[2393]: Server listening on :: port 22.
Jun 21 17:40:43 hh23-12 sshd[12657]: Connection closed by 171.70.163.201 [preauth]
Jun 21 17:41:13 hh23-12 sshd[12659]: Accepted password for root from 171.70.163.201 port 40499
Jun 21 17:46:41 hh23-12 systemd[1]: Stopping OpenSSH server daemon...
Jun 21 17:46:41 hh23-12 sshd[2393]: Received signal 15; terminating.
Jun 21 17:46:41 hh23-12 systemd[1]: Started OpenSSH server daemon.
Jun 21 17:46:41 hh23-12 systemd[1]: Starting OpenSSH server daemon...
Jun 21 17:46:41 hh23-12 sshd[13930]: Server listening on 0.0.0.0 port 22.
Jun 21 17:46:41 hh23-12 sshd[13930]: Server listening on :: port 22.
Jun 21 17:50:45 hh23-12 sshd[33964]: Accepted password for root from 171.70.163.201 port 40545
Jun 21 17:56:36 hh23-12 sshd[34028]: Connection closed by 192.168.212.20 [preauth]
Jun 21 17:57:08 hh23-12 sshd[34030]: Accepted publickey for root from 10.117.212.20 port

```

```
62819
Jun 22 16:42:40 hh23-12 sshd[8485]: Invalid user user1 from 10.117.212.20
Jun 22 16:42:40 hh23-12 sshd[8485]: input_userauth_request: invalid user user1 [preauth]
s
```

Verifying SELinux

To minimize the impact of a security breach on a Cisco NFVI server, the Cisco VM enables SELinux (Security Enhanced Linux) to protect the server resources. To validate that SELinux is configured and running in enforcing mode, use the `sestatus` command to view the status of SELinux and verify that its status is enabled and in enforcing mode. For example:

```
[root@mgmt1 ~]# /usr/sbin/sestatus -v
SELinux status:                enabled
SELinuxfs mount:              /sys/fs/selinux
SELinux root directory:       /etc/selinux
Loaded policy name:            targeted
Current mode:                  enforcing
Mode from config file:         permissive
Policy MLS status:             enabled
Policy deny_unknown status:    allowed
Max kernel policy version:     28
```

Validating Port Listening Services

To prevent access by unauthorized users and processes, Cisco NFVI has no extra services listening on network ports. To verify this, use the `netstat -plnt` command to get a list of all services listening on the node and verify that no unauthorized services are listening. For example:

```
[root@-control-server-1 ~]# netstat -plnt
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State       PID/Program
name
tcp        0      0 23.23.4.101:8776       0.0.0.0:*                 LISTEN      24468/python2
tcp        0      0 23.23.4.101:5000      0.0.0.0:*                 LISTEN      19874/httpd
tcp        0      0 23.23.4.101:5672      0.0.0.0:*                 LISTEN      18878/beam.smp

tcp        0      0 23.23.4.101:3306      0.0.0.0:*                 LISTEN      18337/mysqld
tcp        0      0 127.0.0.1:11211       0.0.0.0:*                 LISTEN      16563/memcached
tcp        0      0 23.23.4.101:11211    0.0.0.0:*                 LISTEN      16563/memcached
tcp        0      0 23.23.4.101:9292      0.0.0.0:*                 LISTEN      21175/python2
tcp        0      0 23.23.4.101:9999      0.0.0.0:*                 LISTEN      28555/python
tcp        0      0 23.23.4.101:80        0.0.0.0:*                 LISTEN      28943/httpd
tcp        0      0 0.0.0.0:4369          0.0.0.0:*                 LISTEN      18897/epmd

tcp        0      0 127.0.0.1:4243       0.0.0.0:*                 LISTEN      14673/docker

tcp        0      0 0.0.0.0:22            0.0.0.0:*                 LISTEN      2909/ssh

tcp        0      0 23.23.4.101:4567      0.0.0.0:*                 LISTEN      18337/mysqld
tcp        0      0 23.23.4.101:15672    0.0.0.0:*                 LISTEN      18878/beam.smp
tcp        0      0 0.0.0.0:35672         0.0.0.0:*                 LISTEN      18878/beam.smp
tcp        0      0 127.0.0.1:25         0.0.0.0:*                 LISTEN      4531/master
tcp        0      0 23.23.4.101:35357    0.0.0.0:*                 LISTEN      19874/httpd
```

```

tcp        0      0 23.23.4.101:8000    0.0.0.0:*        LISTEN     30505/python
tcp        0      0 23.23.4.101:6080    0.0.0.0:*        LISTEN     27996/python2
tcp        0      0 23.23.4.101:9696    0.0.0.0:*        LISTEN     22396/python2
tcp        0      0 23.23.4.101:8004    0.0.0.0:*        LISTEN     30134/python
tcp        0      0 23.23.4.101:8773    0.0.0.0:*        LISTEN     27194/python2
tcp        0      0 23.23.4.101:8774    0.0.0.0:*        LISTEN     27194/python2
tcp        0      0 23.23.4.101:8775    0.0.0.0:*        LISTEN     27194/python2
tcp        0      0 23.23.4.101:9191    0.0.0.0:*        LISTEN     20752/python2
tcp6       0      0 :::9200             :::*             LISTEN     18439/xinetd
tcp6       0      0 :::4369             :::*             LISTEN     18897/epmd
tcp6       0      0 :::22               :::*             LISTEN     2909/sshd
tcp6       0      0 :::1:25             :::*             LISTEN     4531/master

```

Validating Non-Root Users for OpenStack Services

To prevent unauthorized access, Cisco NFVI runs OpenStack processes as a non-root user. To verify OpenStack processes are not running as root, use the `ps` command to get a list of all node processes. In the following example the user is 162:

```

[root@control-server-1 ~]# ps -aux | grep nova-api
162      27194  0.6  0.0 360924 132996 ?        S    Aug08   76:58 /usr/bin/python2
/usr/bin/nova-api
162      27231  0.0  0.0 332192 98988 ?        S    Aug08    0:01 /usr/bin/python2
/usr/bin/nova-api
162      27232  0.0  0.0 332192 98988 ?        S    Aug08    0:01 /usr/bin/python2
/usr/bin/nova-api
162      27233  0.0  0.0 332192 98988 ?        S    Aug08    0:01 /usr/bin/python2
/usr/bin/nova-api

```

Verifying Password Strength

Cisco NFVI passwords can be generated in two ways during installation:

- The Cisco NFVI installer generates unique passwords automatically for each protected service.
- You can provide an input file containing the passwords you prefer.

Cisco-generated passwords are unique, long, and contain a mixture of uppercase, lowercase, and numbers. If you provide the passwords, password strength is your responsibility.

You can view the passwords by displaying the `secrets.yaml` file. For example:

```

[root@mgmt1 ~]# cat ~/openstack-configs/secrets.yaml
ADMIN_USER_PASSWORD: QaZ12n13wvVNY7AH
CINDER_DB_PASSWORD: buJL8pAfytoJ0Icm

```

```
CINDER_KEYSTONE_PASSWORD: AYbcB8mx6a5Ot549
CLOUDPULSE_KEYSTONE_PASSWORD: HAT6vbl7Z56yZLtN
COBBLER_PASSWORD: bax8leYFyyDon0ps
CPULSE_DB_PASSWORD: aYGSzURpGChztbMv
DB_ROOT_PASSWORD: bjb3Uvwus6cvaNe5
KIBANA_PASSWORD: c50e57Dbm7LF0dRV
[...]
```

Reconfiguring Passwords and OpenStack Configurations



Note This topic does not apply if you have installed the optional Cisco Virtual Topology System. For information about use of passwords when VTS is installed, see the *Installing Cisco VTS* section in the *Cisco NFV Infrastructure 2.2 Installation Guide*.

You can reset some configurations after installation including the OpenStack service password and debugs, TLS certificates, and ELK configurations. Two files, `secrets.yaml` and `openstack_config.yaml` which are located in `:/root/installer-{tag id}/openstack-configs/`, contain the passwords, debugs, TLS file location, and ELK configurations. Also, Elasticsearch uses disk space for the data that is sent to it. These files can grow in size, and Cisco VIM has configuration variables that establishes the frequency and file size under which they are rotated.

Cisco VIM installer generates the OpenStack service and database passwords with 16 alphanumeric characters and stores those in `/root/openstack-configs/secrets.yaml`. You can change the OpenStack service and database passwords using the password reconfigure command on the deployed cloud. The command identifies the containers affected by the password change and restarts them so the new password can take effect.



Note Always schedule password reconfiguration in a maintenance window because container restarts might disrupt the control plane

Run the following command to view the list of passwords and configurations :

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 installer-xxxx]# ciscovim help reconfigure
usage: ciscovim reconfigure [--regenerate_secrets] [--setpassword <secretkey>]
                             [--setopenstackconfig <option>]
```

Reconfigure the openstack cloud

Optional arguments:

```
--regenerate_secrets          Regenerate All Secrets
--setpassword <secretkey>     Set of secret keys to be changed.
--setopenstackconfig <option> Set of Openstack config to be changed.
```

```
[root@mgmt1 ~]# ciscovim list-openstack-configs
```

Name	Option
CINDER_DEBUG_LOGGING	False
KEYSTONE_DEBUG_LOGGING	False
CLOUDPULSE_VERBOSE_LOGGING	True
MAGNUM_VERBOSE_LOGGING	True
NOVA_DEBUG_LOGGING	True
NEUTRON_VERBOSE_LOGGING	True

```

|     external_lb_vip_cert | /root/openstack-configs/haproxy.pem |
| GLANCE_VERBOSE_LOGGING | True |
| | |
|     elk_rotation_frequency | monthly |
| CEILOMETER_VERBOSE_LOGGING | True |
|     elk_rotation_del_older | 10 |
| HEAT_DEBUG_LOGGING | False |
| KEYSTONE_VERBOSE_LOGGING | True |
|     external_lb_vip_cacert | /root/openstack-configs/haproxy-ca.crt |
| MAGNUM_DEBUG_LOGGING | True |
| CINDER_VERBOSE_LOGGING | True |
|     elk_rotation_size | 2 |
| CLOUDPULSE_DEBUG_LOGGING | False |
| NEUTRON_DEBUG_LOGGING | True |
| HEAT_VERBOSE_LOGGING | True |
| CEILOMETER_DEBUG_LOGGING | False |
| GLANCE_DEBUG_LOGGING | False |
| NOVA_VERBOSE_LOGGING | True |
+-----+
[root@mgmt1 installer-xxxx]#
[root@mgmt1 installer-xxxx]# ciscovim list-password-keys
+-----+
| Password Keys |
+-----+
| COBBLER_PASSWORD |
| CPULSE_DB_PASSWORD |
| DB_ROOT_PASSWORD |
| KIBANA_PASSWORD |
| GLANCE_DB_PASSWORD |
| GLANCE_KEYSTONE_PASSWORD |
| HAPROXY_PASSWORD |
| HEAT_DB_PASSWORD |
| HEAT_KEYSTONE_PASSWORD |
| HEAT_STACK_DOMAIN_ADMIN_PASSWORD |
| HORIZON_SECRET_KEY |
| KEYSTONE_ADMIN_TOKEN |
| KEYSTONE_DB_PASSWORD |
| METADATA_PROXY_SHARED_SECRET |
| NEUTRON_DB_PASSWORD |
| NEUTRON_KEYSTONE_PASSWORD |
| NOVA_DB_PASSWORD |
| NOVA_KEYSTONE_PASSWORD |
| RABBITMQ_ERLANG_COOKIE |
| RABBITMQ_PASSWORD |
| WSREP_PASSWORD |
+-----+
[root@mgmt1 installer-xxxx]#

```

You can change specific password and configuration identified from the available list.

Run the reconfiguration command as follows:

```

[root@mgmt1 ~]# ciscovim help reconfigure
usage: ciscovim reconfigure [--regenerate_secrets] [--setpassword <secretkey>]
                             [--setopenstackconfig <option>]

```

Reconfigure the Openstack cloud

Optional arguments:

```

--regenerate_secrets          Regenerate All Secrets
--setpassword <secretkey>     Set of secret keys to be changed.
--setopenstackconfig <option> Set of Openstack config to be changed.

```

```

[root@mgmt1 ~]# ciscovim reconfigure --setpassword ADMIN_USER_PASSWORD,NOVA_DB_PASSWORD
--setopenstackconfig HEAT_DEBUG_LOGGING,HEAT_VERBOSE_LOGGING

```

```

Password for ADMIN_USER_PASSWORD:
Password for NOVA_DB_PASSWORD:
Enter T/F for option HEAT_DEBUG_LOGGING:T
Enter T/F for option HEAT_VERBOSE_LOGGING:T

```

The password must be alphanumeric and can be maximum 32 characters in length.

Following are the configuration parameters for OpenStack:

Configuration Parameter	Allowed Values
CEILOMETER_DEBUG_LOGGING	T/F (True or False)
CEILOMETER_VERBOSE_LOGGING	T/F (True or False)
CINDER_DEBUG_LOGGING	T/F (True or False)
CINDER_VERBOSE_LOGGING	T/F (True or False)
CLOUDPULSE_DEBUG_LOGGING	T/F (True or False)
CLOUDPULSE_VERBOSE_LOGGING	T/F (True or False)
GLANCE_DEBUG_LOGGING	T/F (True or False)
GLANCE_VERBOSE_LOGGING	T/F (True or False)
HEAT_DEBUG_LOGGING	T/F (True or False)
HEAT_VERBOSE_LOGGING	T/F (True or False)
KEYSTONE_DEBUG_LOGGING	T/F (True or False)
KEYSTONE_VERBOSE_LOGGING	T/F (True or False)
MAGNUM_DEBUG_LOGGING	T/F (True or False)
MAGNUM_VERBOSE_LOGGING	T/F (True or False)
NEUTRON_DEBUG_LOGGING	T/F (True or False)
NEUTRON_VERBOSE_LOGGING	T/F (True or False)
NOVA_DEBUG_LOGGING	T/F (True or False)
NOVA_VERBOSE_LOGGING	T/F (True or False)
elk_rotation_del_older	Days after which older logs are purged
elk_rotation_frequency	Available options: "daily", "weekly", "fortnightly", "monthly"
elk_rotation_size	Gigabytes (entry of type float/int is allowed)
external_lb_vip_cacert	Location of HAProxy CA certificate
external_lb_vip_cert	Location of HAProxy certificate

NOVA_RAM_ALLOCATION_RATIO	Mem oversubscription ratio (from 1.0 to 4.0)
NOVA_CPU_ALLOCATION_RATIO	CPU allocation ratio (from 1.0 to 16.0)
ES_SNAPSHOT_AUTODELETE	Elastic search auto-delete configuration, can manage the following: period: ["hourly", "daily", "weekly", "monthly"] # Frequency of cronjob to check for disk space threshold_warning: <1-99> # % of disk space occupied to display warning message threshold_low: <1-99> # % of disk space occupied after cleaning up snapshots threshold_high: <1-99> # % of disk space when starting to delete snapshots

Alternatively, you can regenerate all passwords using `regenerate_secrets` command option as follows:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim reconfigure --regenerate_secrets
```

In addition to the services passwords, you can change the debug and verbose options for Heat, Glance, Cinder, Nova, Neutron, Keystone and Cloudpulse in `/root/openstack-configs/openstack_config.yaml`. You can modify the other configurations including the ELK configuration parameters, API and Horizon TLS certificates, Root CA, `NOVA_EAMALLOCATION_RATIO` and `ES_SNAPSHOT_AUTODELETE`. When reconfiguring these options (For Example API and TLS), some control plane downtime will occur, so plan the changes during maintenance windows.

The command to reconfigure these elements are:

```
ciscovim reconfigure
```

The command includes a built-in validation to ensure you do not enter typos in the `secrets.yaml` or `openstack_config.yaml` files.

When reconfiguration of password or enabling of `openstack-services` fails, all subsequent pod management operations are blocked. In such cases, we recommend that you contact Cisco TAC to resolve the situation.

Enabling NFVIMON Post Pod Install

The dispatcher is the only component in NFVIMON offering that is managed by VIM orchestrator. While the dispatcher acts as a conduit to pass openstack information of the pod to the collectors, it is the Cisco NFVI Zenpack sitting in the CC/RM node, that gathers the node level information. To enable dispatcher as part of the VIM Install, update the `setup_data` with the following information:

```
#Define the PODNAME
PODNAME: <PODNAME with no space>; ensure that this is unique across all the pods
NFVIMON:
  MASTER:
    # Master Section
    admin_ip: <IP address of Control Centre VM>
  COLLECTOR:
    # Collector Section
management_vip: <VIP for ceilometer/dispatcher to use> #Should be unique across the VIM
Pod; Should be part of br_mgmt network
Collector_VM_Info:
```

```

-
  hostname: <hostname of Collector VM 1>
  password: <password_for_collector_vm1> # max length of 32
  ccuser_password: <password from master for 'ccuser' (to be used for self monitoring)>
# max length of 32
  admin_ip: <ssh_ip_collector_vm1> # Should be part of br_api network
  management_ip: <mgmt_ip_collector_vm1> # Should be part of br_mgmt network
-
  hostname: <hostname of Collector VM 2>
  password: <password_for_collector_vm2> # max length of 32
  ccuser_password: <password from master for 'ccuser' (to be used for self monitoring)>
# max length of 32
  admin_ip: <ssh_ip_collector_vm2> # Should be part of br_api network
  management_ip: <mgmt_ip_collector_vm2> # Should be part of br_mgmt network
DISPATCHER:
  rabbitmq_username: admin # Pod specific user for dispatcher module in
ceilometer-collector

```

To monitor TOR, ensure that the following TORSWITCHINFO sections are defined in the setup_data.yaml.

```

TORSWITCHINFO:
  SWITCHDETAILS:
-
  hostname: <switch_a_hostname>: # Mandatory for NFVIMON if switch monitoring is
needed
  username: <TOR switch username> # Mandatory for NFVIMON if switch monitoring is
needed
  password: <TOR switch password> # Mandatory for NFVBENCH; Mandatory for NFVIMON
if switch monitoring is needed
  ssh_ip: <TOR switch ssh ip> # Mandatory for NFVIMON if switch monitoring is
needed
  ....
-
  hostname: <switch_b_hostname>: # Mandatory for NFVIMON if switch monitoring is
needed
  username: <TOR switch username> # Mandatory for NFVIMON if switch monitoring is
needed
  password: <TOR switch password> # Mandatory for NFVIMON if switch monitoring is
needed
  ssh_ip: <TOR switch ssh ip> # Mandatory for NFVIMON if switch monitoring is
needed
  ....

```

To initiate the integration of NFVIMON on an existing pod, copy the setupdata into a local dir and update it manually with information listed above, and then run reconfiguration command as follows:

```

[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include NFVIMON related
info)
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure

```

It should be noted that un-configuration of this feature is not supported today. Additionally, NFVIMON is supported only on a pod running with Keystone v2.

Reconfiguring CIMC Password on an Existing Install

Cisco VIM, allows you to reconfigure the CIMC password on an existing install along with OpenStack services.



Note You must have a C-series pod, up and running with Cisco to reconfigure the CIMC password.

Step 1 Update the `cimc_password` in the CIMC-COMMON section, and/or the individual `cimc_password` for each server and then run the reconfigure option provided by `Ciscovimclient`.

```
CIMC-COMMON:
  cimc_username: "admin"
  cimc_password: <"new password">
:
:
SERVERS:
:
control-server-2:
  cimc_info: {'cimc_ip': '<ip_addr>',
             'cimc_username': 'admin',
             'cimc_password': '<update with new password>'} # only needed if each server has specific
password
```

Step 2 To change the CIMC password for the pod, copy the `setupdata` into a local location and update it manually with the CIMC password as shown in the snippet above. The new password must satisfy atleast three of the following conditions:

Note Do not change CIMC password directly into the exiting `/root/openstack-configs/setup_data.yaml` file.

- Must contain at least one lower case letter.
- Must contain at least one upper case letter.
- Must contain at least one digit between 0 to 9.
- One of these special characters `!$#@%^_+=*&`
- Your password has to be 8 to 14 characters long.

Step 3 Run the vim reconfiguration command to post update the `setup_data` as follows:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
[root@mgmt1 ~]# cp <my_setup_data.yaml> <my_setup_data_original.yaml>
[root@mgmt1 ~]# vi my_setup_data.yaml (update the relevant CIMC setup_data to include LDAP info)
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

Note After successful completion of the CIMC Password, reconfigure operation triggers an auto-back when the management node auto-back recovery feature is enabled. If the CIMC Password reconfigure fails, contact Cisco TAC to recover from the failure.

Increasing Provider and Tenant VLAN Ranges

Cisco VIM, provides the flexibility of increasing the provider and tenant VLAN ranges after the post pod installation. Increasing provider and tenant VLAN ranges applies to C-series and B-series pod that is enabled with Cisco UCS Manager plugin. B-series pod running without Cisco UCS Manager plugin, cannot use this feature because of the inherent day-0 networking configuration to be done in FI.



Note You should have the tenant and provider networks enabled on the pod from day-0.

To increase provider and tenant VLAN ranges enter the `TENANT_VLAN_RANGES` and/or `PROVIDER_VLAN_RANGES` in the `setup_data.yaml` file and run the reconfigure command through `Ciscovimclient` as follows:

```
TENANT_VLAN_RANGES: old_vlan_info, new_vlan_info
or/and
PROVIDER_VLAN_RANGES: old_vlan_info, new_vlan_info
```

To change the pod, copy the setupdata into a local dir and update it manually by running the following command:

```
[root@mgmt1 ~]# cd /root/ [root@mgmt1 ~]# mkdir MyDir [root@mgmt1 ~]# cd MyDir
```

Update the `setup_data`, by running the following command:

```
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml> [root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data with the right info)
```

Run the re-configuration command as follows:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ./ciscovimclient/ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

Fernet Key Operations

Keystone fernet token format is based on the cryptographic authentication method - Fernet. Fernet is an implementation of Symmetric Key Encryption. Symmetric key encryption is a cryptographic mechanism that uses the same cryptographic key to encrypt plaintext and the same cryptographic key to decrypt ciphertext. Fernet authentication method also supports multiple keys where it takes a list of symmetric keys, performs all encryption using the first key in a list and attempts to decrypt using all the keys from that list.

The Cisco NFVI pods uses Fernet keys by default. The following operations can be carried out in Cisco NFVI pods.

To check if the fernet keys are successfully synchronized across the keystone nodes.

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim help check-fernet-keys
usage: ciscovim check-fernet-keys
```

Check whether the fernet keys are successfully synchronized across keystone nodes.

To set the fernet key frequency:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim help period-rotate-fernet-keys
usage: ciscovim period-rotate-fernet-keys <SET_PERIOD_ROTATION_FERNET_KEYS>
```

```
Set the frequency of fernet keys rotation on keystone
Positional arguments:
  <SET_PERIOD_ROTATION_FERNET_KEYS>
Frequency to set for period rotation
```

To forcefully rotate the fernet keys:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim help rotate-fernet-keys
usage: ciscovim rotate-fernet-keys
Trigger rotation of the fernet keys on keystone
```

To resync the fernet keys across the keystone nodes:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim help resync-fernet-keys
usage: ciscovim resync-fernet-keys
Resynchronize the fernet keys across all the keystone nodes
```

Managing Certificates

When TLS protection is configured for the OpenStack APIs, the two certificate files, `haproxy.pem` and `haproxy-ca.crt`, are stored in the `/root/openstack-configs/` directory. Clients running on servers outside of the deployed cloud to verify cloud authenticity need a copy of the root certificate (`haproxy-ca.crt`). If a well-known certificate authority has signed the installed certificate, no additional configuration is needed on client servers. However, if a self-signed or local CA is used, copy `haproxy-ca.crt` to each client. Following instructions specific to the client operating system or browser to install the certificate as a trusted certificate.

Alternatively, you can explicitly reference the certificate when using the OpenStack CLI by using the environment variable `OS_CACERT` or command line parameter `-cacert`.

While Cisco NFVI is operational, a daily check is made to monitor the expiration dates of the installed certificates. If certificates are not nearing expiration, an informational message is logged. As the certificate approaches expiration, an appropriate warning or critical message is logged.

```
2017-04-24T13:56:01 INFO Certificate for OpenStack Endpoints at 192.168.0.2:5000 expires
in 500 days
```

It is important to replace the certificates before they expire. After Cisco NFVI is installed, you can update the certificates by replacing the `haproxy.pem` and `haproxy-ca.crt` files and running the `reconfigure` command:

```
cd ~/installer-xxxx; ciscovim reconfigure
```

Reconfiguring TLS Certificates

Cisco VIM provides a way to configure TLS certificates on-demand for any reason. For Example: certificate expiration policies governing certificate management.

Reconfiguration of certificates in general is supported in the following components:

- Cisco VIM Rest API endpoints:

Steps to be performed to reconfigure certificate files are as follows:

- Copy the new key, CA root and certificate files into the `~/openstack-configs` folder under the following filenames

```
cp <new-ca-root-cert> ~/openstack-configs/mercury-ca.crt
cp <new-key-file> ~/openstack-configs/mercury.key
cp <new-cert-file> ~/openstack-configs/mercury.crt
```

- Once copied run the reconfigure steps as under:

```
cd ~/installer-xxxx/tools
./restapi.py -a reconfigure-tls
```

- OpenStack API endpoints

Steps to be performed to reconfigure certificate files are as follows:

- Copy the new key, CA root and certificate files into the ~/openstack-configs folder under the following filenames

```
cp <new-ca-root-cert> ~/openstack-configs/haproxy-ca.crt
cp <new-cert-file> ~/openstack-configs/haproxy.pem
```

- Once copied run the reconfigure steps as follows:

```
cd ~/installer-xxxx; ciscovim reconfigure
```

- SwiftStack Service through Horizon and CinderBackup Service.

- Reconfiguring TLS certificates for SwiftStack mainly involves client side certificate updates. The CA root certificate in both these cases is updated for components within OpenStack that are clients of the SwiftStack service in general.
- Copy the new CA root certificate to the ~/openstack-configs folder and run reconfigure.

```
cp <new-ca-root-cert> ~/openstack-configs/haproxy-ca.crt
cd ~/installer-xxxx; ciscovim reconfigure
```

- Logstash service and Fluentd (client-side certificates).

- For the Logstash service on the management node, both the key and certificate file are reconfigured as part of the reconfigure operation.
- For the Fluentd service on the controllers, compute and storage nodes, the certificate file are reconfigured as part of the reconfigure operation.
- Copy of the key and certificate files to the ~/openstack-configs folder on the management node and run reconfigure operation.

```
cp <new-key-file> ~/openstack-configs/logstash-forwarder.key
cp <new-cert-file> ~/openstack-configs/logstash-forwarder.crt
cd ~/installer-xxxx; ciscovim reconfigure
```

Enabling Keystone v3 on an Existing Install

To continue enhancing our security portfolio, and multi-tenancy with the use of domains, Keystone v3 support has been added in Cisco VIM from an authentication end-point. It should be noted that Keystone v2 and v3 are mutually exclusive. The administrator has to decide during install time the authentication end-point version to go with. By default, VIM orchestrator picks keystone v2 as the authentication end-point. So one can enable Keystonev3 as an install option on day-0 (see 2.2 CiscoVIM install guide), or enable it as a reconfigure option

after the pod is installed. To enable Keystone v3 after the pod is installed, one needs to define the following under the optional service section in the `setup_data.yaml` file.

```
# Optional Services:
OPTIONAL_SERVICE_LIST:
- keystonev3
```

To initiate the integration of Keystone v3 on an existing pod, copy the `setupdata` into a local dir and update it manually, then run reconfiguration command as follows:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include keystone v3 info)
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

It should be noted that un-configuration of this feature is not supported today. Additionally, re-versioning Keystone API from v3 to v2 is also not supported.

LDAP/AD support with Keystone v3

With the introduction of KeystoneV3, the openstack service authentication can now be delegated to an external LDAP/AD server. In Cisco VIM, this feature has been introduced optionally if the authorization is done by Keystone v3. Just like Keystonev3, this feature can be enabled on an existing pod running Cisco VIM. To avail of this feature post pod deployment, the `setup_data` needs to be augmented with the following information during the installation of the pod.

An important pre-requisite for enabling AD/LDAP integration is that the AD/LDAP endpoint **MUST** be reachable from all the Controller nodes that run OpenStack Keystone Identity Service.

```
LDAP:
  domain: <Domain specific name>
  user_objectclass: <objectClass for Users> # e.g organizationalPerson
  group_objectclass: <objectClass for Groups> # e.g. groupOfNames
  user_tree_dn: '<DN tree for Users>' # e.g. 'ou=Users,dc=cisco,dc=com'
  group_tree_dn: '<DN tree for Groups>' # e.g. 'ou=Groups,dc=cisco,dc=com'
  suffix: '<suffix for DN>' # e.g. 'dc=cisco,dc=com'
  url: '<ldap:// host:port>' # e.g. 'ldap://172.26.233.104:389'
or
url: '<ldaps|ldap>://[<ip6-address>]:[port] '
e.g.ldap://[2001:420:293:2487:d1ca:67dc:94b1:7e6c]:389 ---> note the mandatory "[.. ]"
around the ipv6 address
  user: '<DN of bind user>' # e.g. 'dc=admin,dc=cisco,dc=com'
  password: <password> # e.g. password of bind user

user_filter = (memberOf=CN=os-users,OU=OS-Groups,DC=mercury,DC=local)
user_id_attribute = sAMAccountName
user_name_attribute = sAMAccountName
user_mail_attribute = mail # Optional
group_tree_dn = ou=OS-Groups,dc=mercury,dc=local
group_name_attribute = sAMAccountName
```

To initiate the integration of LDAP with Keystone v3 on an existing pod, copy the `setupdata` into a local dir and update it manually with the relevant LDAP and Keystone v3 (if absent from before) configuration, then run reconfiguration command as follows:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
```

```
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include LDAP info)
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

The reconfigure feature supports a full or partial reconfiguration of the LDAP integration service.



Note All the parameters within the LDAP stanza are configurable with the exception of the domain parameter.

Integrating identity with LDAP over TLS: The automation supports keystone integration with LDAP over TLS. In order to enable TLS, the CA root certificate must be presented as part of the /root/openstack-configs/haproxy-ca.crt file. The url parameter within the LDAP stanza must be set to ldaps.

Additionally, the url parameter supports the following format: url: '<ldaps | ldap>://<FQDN | IP-Address>:[port]'

The protocol can be one of the following: ldap for non-ssl and ldaps when TLS has to be enabled.

The ldap host can be a fully-qualified domain name (FQDN) or an IPv4 or v6 Address depending on how the SSL certificates are generated. .

The port number is optional and if not provided assumes that the ldap services are running on the default ports. For Example: . 389 for non-ssl and 636 for ssl. However, if these are not the defaults, then the non-standard port numbers must be provided. Except for the domain, all other item values can be changed via the 'reconfigure' option.



CHAPTER 5

Managing Cisco NFVI Storage

This chapter describes basic architectural concepts that will help you understand the Cisco NFVI data storage architecture and data flow. It also provides techniques you can use to monitor the storage cluster health and the health of all systems that depend on it

- [Cisco NFVI Storage Architecture, on page 115](#)
- [Verifying and Displaying Ceph Storage Pools, on page 116](#)
- [Checking the Storage Cluster Health, on page 117](#)
- [Checking Glance Connectivity, on page 118](#)
- [Verifying Glance and Ceph Monitor Keyrings, on page 119](#)
- [Verifying Glance Image ID on Ceph, on page 120](#)
- [Checking Cinder Connectivity, on page 120](#)
- [Verifying the Cinder and Ceph Monitor Keyrings, on page 121](#)
- [Verifying the Cinder Volume ID on Ceph, on page 122](#)
- [Checking Nova Connectivity, on page 122](#)
- [Verifying the Nova and Ceph Monitor Keyrings, on page 123](#)
- [Verifying Nova Instance ID, on page 124](#)
- [Displaying Docker Disk Space Usage, on page 125](#)
- [Reconfiguring SwiftStack Integration, on page 125](#)
- [Reconfiguring Administrator Source Networks, on page 127](#)
- [Password Reset for Cisco VIM Management Node, on page 128](#)

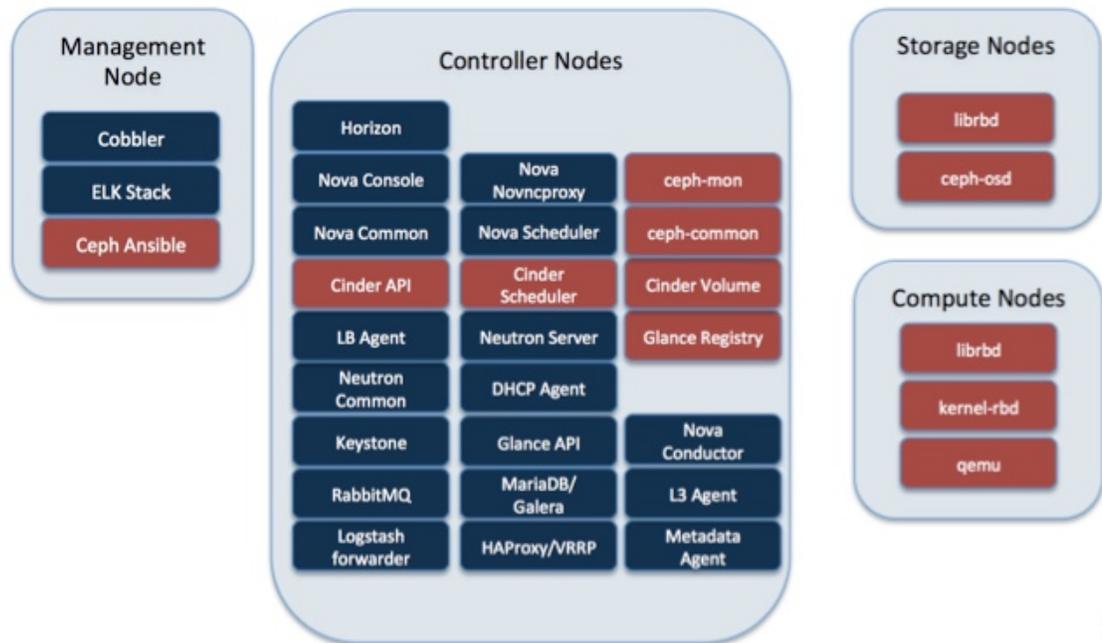
Cisco NFVI Storage Architecture

OpenStack has multiple storage back ends. Cisco NFVI uses the Ceph back end. Ceph supports both block and object storage and is therefore used to store VM images and volumes that can be attached to VMs. Multiple OpenStack services that depend on the storage backend include:

- Glance (OpenStack image service)—Uses Ceph to store images.
- Cinder (OpenStack storage service)—Uses Ceph to create volumes that can be attached to VMs.
- Nova (OpenStack compute service)—Uses Ceph to connect to the volumes created by Cinder.

The following figure shows the Cisco NFVI storage architecture component model.

Figure 18: Cisco NFVI Storage Architecture



41-1023

Verifying and Displaying Ceph Storage Pools

Ceph is configured with four independent pools: images, volumes, vms, and backups. (A default rbd pool is used internally.) Each Ceph pool is mapped to an OpenStack service. The Glance service stores data in the images pool, and the Cinder service stores data in the volumes pool. The Nova service can use the vms pool to boot ephemeral disks directly from the Ceph cluster depending on how the NOVA_BOOT_FROM option in the `~/openstack-configs/setup_data.yaml` was configured prior to Cisco NFVI installation. If NOVA_BOOT_FROM is set to ceph before you run the Cisco NFVI installation, the Nova service boot up from the Ceph vms pool. By default, NOVA_BOOT_FROM is set to local, which means that all VM ephemeral disks are stored as files in the compute nodes. Changing this option after installation does not affect the use of the vms pool for ephemeral disks.

The Glance, Cinder, and Nova OpenStack services depend on the Ceph cluster for backend storage. Therefore, they need IP connectivity to the controller nodes. The default port used to connect Glance, Cinder, and Nova to the Ceph cluster is 6789. Authentication through cephx is required, which means authentication tokens, called keyrings, must be deployed to the OpenStack components for authentication.

To verify and display the Cisco NFVI Ceph storage pools:

Step 1 Launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 3 List the Ceph pools:

```
cephmon_4612 [root@controller_server-1 ~]# ceph osd lspools
0 rbd,1 images,2 volumes,3 vms,4 backups,
```

Step 4 List the images pool content:

```
cephmon_4612 [ceph@controller_server-1 /]$ rbd list images
a4963d51-d3b7-4b17-bf1e-2ebac07e1593
```

Checking the Storage Cluster Health

Cisco recommends that you perform a few verifications to determine whether the Ceph cluster is healthy and is connected to the Glance, Cinder, and Nova OpenStack services, which have Ceph cluster dependencies. The first task to check the health of the cluster itself by completing the following steps:

Step 1 From the Cisco NFVI management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 3 Check the Ceph cluster status:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ ceph status
```

Sample response:

```
cluster dbc29438-d3e0-4e0c-852b-170aaf4bd935
health HEALTH_OK
monmap e1: 3 mons at {ceph-controller_server-1=20.0.0.7:6789/0,
ceph-controller_server-2=20.0.0.6:6789/0,ceph-controller_server-3=20.0.0.5:6789/0}
election epoch 8, quorum 0,1,2 ceph-controller_server-3,
ceph-controller_server-2,ceph-controller_server-1
osdmap e252: 25 osds: 25 up, 25 in
pgmap v593: 1024 pgs, 5 pools, 406 MB data, 57 objects
2341 MB used, 61525 GB / 61527 GB avail
1024 active+clean
```

This example displays three monitors, all in good health, and 25 object storage devices (OSDs). All OSDs show as up and in the cluster.

Step 4 To see a full listing of all OSDs sorted by storage node, enter:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ ceph osd tree
```

Sample response:

ID	WEIGHT	TYPE	NAME	UP/DOWN	REWEIGHT	PRIMARY-AFFINITY
-1	60.18979	root	default			
-2	18.96994	host	controller_server-2			
1	2.70999	osd	osd.1	up	1.00000	1.00000
5	2.70999	osd	osd.5	up	1.00000	1.00000
6	2.70999	osd	osd.6	up	1.00000	1.00000

```

11 2.70999      osd.11      up 1.00000      1.00000
12 2.70999      osd.12      up 1.00000      1.00000
17 2.70999      osd.17      up 1.00000      1.00000
20 2.70999      osd.20      up 1.00000      1.00000
-3 18.96994     host controller_server-1
 0 2.70999      osd.0       up 1.00000      1.00000
 4 2.70999      osd.4       up 1.00000      1.00000
 8 2.70999      osd.8       up 1.00000      1.00000
10 2.70999      osd.10      up 1.00000      1.00000
13 2.70999      osd.13      up 1.00000      1.00000
16 2.70999      osd.16      up 1.00000      1.00000
18 2.70999      osd.18      up 1.00000      1.00000
-4 18.96994     host controller_server-3
 2 2.70999      osd.2       up 1.00000      1.00000
 3 2.70999      osd.3       up 1.00000      1.00000
 7 2.70999      osd.7       up 1.00000      1.00000
 9 2.70999      osd.9       up 1.00000      1.00000
14 2.70999      osd.14      up 1.00000      1.00000
15 2.70999      osd.15      up 1.00000      1.00000
19 2.70999      osd.19      up 1.00000      1.00000
-5 3.27997     host controller_server-4
21 0.81999      osd.21      up 1.00000      1.00000
22 0.81999      osd.22      up 1.00000      1.00000
23 0.81999      osd.23      up 1.00000      1.00000
24 0.81999      osd.24      up 1.00000      1.00000

```

What to do next

After you verify the Ceph cluster is in good health, check that the individual OpenStack components have connectivity and their authentication tokens—keyrings—match the Ceph Monitor keyrings. The following procedures show how to check the connectivity and authentication between Ceph and Glance, Ceph and Cinder, and Ceph and Nova.

Checking Glance Connectivity

The Glance API container must be connected to the Cisco NFVI controller nodes. Complete the following steps to verify the Glance to controller node connectivity:

Step 1 From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Navigate to the Glance API container:

```
[root@controller_server-1 ~]# glanceapi
```

Step 3 Check the Glance API container connectivity to a controller node different from the one entered in Step 1, in this case, controller_server 2:

```
glanceapi_4612 [glance@controller_server-1 /]$ curl controller_server-2:6789
```

If the connection is successful, you see a message like the following:

```
glanceapi_4612 [glance@controller_server-1 /]$ curl controller_server-2:6789
ceph v027?
```

If the connection is not successful, you see a message like the following:

```
glanceapi_4612 [glance@controller_server-1 /]$ curl controller_server-2:6789
curl: (7) Failed connect to controller_server-2:6789; Connection refused
```

A message like the one above means the Ceph monitor running on the target controller node `controller_server-2` is not listening on the specified port or there is no route to it from the Glance API container.

Checking one controller node should be enough to ensure one connection path available for the Glance API. However, because Cisco NFVI controller nodes run as part of an HA cluster, you should run Step 3 above targeting all the controller nodes in the Cisco NFVI pod.

What to do next

After you verify the Glance API connectivity to all Cisco NFVI controller nodes, check the Glance keyring to ensure it matches the Ceph monitor keyring.

Verifying Glance and Ceph Monitor Keyrings

Complete the following steps to verify the Glance API keyring matches the Ceph Monitor keyring.

Step 1 Launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Navigate to the Glance API container:

```
[root@controller_server-1 ~]# glanceapi
```

Step 3 Check the Glance keyring content, for example:

```
glanceapi_4612 [glance@controller_server-1 /]$ cat /etc/ceph/client.glance.keyring
[client.glance]
key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Step 4 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 5 Display the Ceph Monitor keyring content:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ cat /etc/ceph/ceph.client.glance.keyring
[client.glance]
```

```
key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Verify the keyring matches the Glance API keyring displayed in Step 3.

What to do next

A final check to ensure that Ceph and Glance are connected is to actually import a Glance image using Horizon or the Glance CLI. After you import an image, compare the IDs seen by Glance and by Ceph. They should match, indicating Ceph is handling the backend for Glance.

Verifying Glance Image ID on Ceph

The following steps verify Ceph is properly handling new Glance images by checking that the image ID for a new Glance image is the same as the image ID displayed in Ceph.

-
- Step 1** From the management node, load the OpenStack authentication variables:
- ```
[root@management-server-cisco ~]# source ~/openstack-configs/openrc
```
- Step 2** Import any Glance image. In the example below, a RHEL 7.1 qcow2 image is used.
- ```
[root@management-server-cisco images]# glance image-create
--name "rhel" --disk-format qcow2 --container-format bare --file
rhel-guest-image-7.1-20150224.0.x86_64.qcow2
```
- Step 3** List the Glance images:
- ```
[root@management-server-cisco images]# glance image-list | grep rhel
| a4963d51-d3b7-4b17-bf1e-2ebac07e1593 | rhel
```
- Step 4** Navigate to the Ceph Monitor container:
- ```
[root@controller_server-1 ~]# cephmon
```
- Step 5** Display the contents of the Ceph images pool:
- ```
cephmon_4612 [ceph@controller_server-1 ceph]$ rbd list images | grep
a4963d51-d3b7-4b17-bf1e-2ebac07e1593
a4963d51-d3b7-4b17-bf1e-2ebac07e1593
```
- Step 6** Verify that the Glance image ID displayed in Step 3 matches the image ID displayed by Ceph.
- 

## Checking Cinder Connectivity

The Cinder volume container must have connectivity to the Cisco NFVI controller nodes. Complete the following steps to verify Cinder volume has connectivity to the controller nodes:

- 
- Step 1** From the management node, launch a SSH session to a controller node, for example:
- ```
[root@management-server-cisco ~]# ssh root@controller_server-1
```
- Step 2** Navigate to the Cinder volume container:
- ```
[root@controller_server-1 ~]# cindervolume
```
- Step 3** Check the Cinder volume container connectivity to a controller node different from the one entered in Step 1, in this case, controller\_server-2:
- ```
cindervolume_4612 [cinder@controller_server-1 /]$ curl controller_server-2:6789
```
- If the connection is successful, you see a message like the following:

```
cindervolume_4612 [cinder@controller_server-1 /]$ curl controller_server-2:6789
ceph v027?
```

If the connection is not successful, you see a message like the following:

```
cindervolume_4612 [cinder@controller_server-1 /]$ curl controller_server-2:6789
curl: (7) Failed connect to controller_server-2:6789; Connection refused
```

A message like the one above means the Ceph monitor running on the target controller node controller_server-2 is not listening on the specified port or there is no route to it from the Cinder volume container.

Checking one controller node should be enough to ensure one connection path is available for the Cinder volume. However, because Cisco NFVI controller nodes run as part of an HA cluster, repeat Step 3 targeting all the controller nodes in the Cisco NFVI pod.

What to do next

After you verify the Cinder volume connectivity to all Cisco NFVI controller nodes, check the Cinder keyring to ensure it matches the Ceph monitor keyring.

Verifying the Cinder and Ceph Monitor Keyrings

Complete the following steps to verify the Cinder volume keyring matches the Ceph Monitor keyring.

Step 1 From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Navigate to the Cinder volume container:

```
[root@controller_server-1 ~]# cindervolume
```

Step 3 Check the Cinder keyring content, for example:

```
cindervolume_4612 [cinder@controller_server-1 /]$ cat /etc/ceph/client.cinder.keyring
[client.cinder]
key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Step 4 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 5 Display the Ceph Monitor keyring content:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ cat /etc/ceph/ceph.client.cinder.keyring
[client.cinder]

key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Verify the keyring matches the Cinder volume keyring displayed in Step 3.

What to do next

As a final Ceph and Cinder connectivity verification, import a Cinder image using Horizon or the Cinder CLI. After you import the image, compare the IDs seen by Cinder and by Ceph. They should match, indicating Ceph is handling the backend for Cinder.

Verifying the Cinder Volume ID on Ceph

The following steps verify Ceph is properly handling new Cinder volumes by checking that the volume ID for a new Cinder volume is the same as the volume ID displayed in Ceph.

Step 1 From the management node, load the OpenStack authentication variables:

```
[root@management-server-cisco ~]# source ~/openstack-configs/openrc
```

Step 2 Create an empty volume:

```
[root@management-server-cisco ~]# cinder create --name ciscovoll 5
```

The preceding command creates a new 5 GB Cinder volume named ciscovoll.

Step 3 List the Cinder volumes:

```
[[root@management-server-cisco ~]# cinder list
+-----+-----+-----+-----+
|                ID                | Status | Migration Status |...
+-----+-----+-----+-----+
| dd188a5d-f822-4769-8a57-c16694841a23 | in-use | -                 |...
```

Step 4 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 5 Display the contents of the Ceph volumes pool:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ rbd list volumes
volume-dd188a5d-f822-4769-8a57-c16694841a23
```

Step 6 Verify that the Cinder volume ID displayed in Step 3 matches the volume ID displayed by Ceph, excluding the "volume-" prefix.

Checking Nova Connectivity

The Nova libvirt container must have connectivity to the Cisco NFVI controller nodes. Complete the following steps to verify Nova has connectivity to the controller nodes:

Step 1 From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@Computenode_server-1
```

Step 2 Navigate to the Nova libvirt container:

```
[root@compute_server-1 ~]# libvirt
```

Step 3 Check the Nova libvirt container connectivity to a controller node, in this case, controller_server 1:

```
novalibvirt_4612 [root@compute_server-1 /]$ curl controller_server-2:6789
```

If the connection is successful, you see a message like the following:

```
novalibvirt_4612 [root@compute_server-1 /]$ curl controller_server-1:6789
ceph v027?
```

If the connection is not successful, you see a message like the following:

```
novalibvirt_4612 [root@compute_server-1 /]$ curl controller_server-1:6789
curl: (7) Failed connect to controller_server-1:6789; Connection refused
```

A message like the one above means the Ceph monitor running on the target controller node controller_server-1 is not listening on the specified port or there is no route to it from the Nova libvirt container.

Checking one controller node should be enough to ensure one connection path available for the Nova libvirt. However, because Cisco NFVI controller nodes run as part of an HA cluster, you should run Step 3 above targeting all the controller nodes in the Cisco NFVI pod.

What to do next

After you verify the Nova libvirt connectivity to all Cisco NFVI controller nodes, check the Nova keyring to ensure it matches the Ceph monitor keyring.

Verifying the Nova and Ceph Monitor Keyrings

Complete the following steps to verify the Nova libvirt keyring matches the Ceph Monitor keyring.

Step 1 From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Navigate to the Nova libvirt container:

```
[root@compute_server-1 ~]# libvirt
```

Step 3 Extract the libvirt secret that contains the Nova libvirt keyring:

```
novalibvirt_4612 [root@compute_server-1 /]# virsh secret-list
UUID                               Usage ...
-----
b5769938-e09f-47cb-bdb6-25b15b557e84  ceph client.cinder ...
```

Step 4 Get the keyring from the libvirt secret:

```
novalibvirt_4612 [root@controller_server-1 /]# virsh secret-get-value
b5769938-e09f-47cb-bdb6-25b15b557e84
AQBAPY1XQCBEBEAARoXvmiwmlSMEyEoXKl/sQA==
```

Step 5 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 6 Display the Ceph Monitor keyring content:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ cat /etc/ceph/ceph.client.cinder.keyring
[client.cinder]
```

```
key = AQBAPY1XQCBEBEBAARoXvmiwmlSMEyEoXK1/sQA==
```

Verify the keyring matches the Nova libvirt keyring displayed in Step 3. Notice that in the above example the Cinder keyring is checked even though this procedure is for the Nova libvirt keyring. This occurs because the Nova services need access to the Cinder volumes and so authentication to Ceph uses the Cinder keyring.

What to do next

Complete a final check to ensure that Ceph and Nova are connected by attaching a Nova volume using Horizon or the Nova CLI. After you attach the Nova volume, check the libvirt domain.

Verifying Nova Instance ID

From the management node, complete the following steps to verify the Nova instance ID:

Step 1 Load the OpenStack authentication variables:

```
[root@management-server-cisco installer]# source ~/openstack-configs/openrc
```

Step 2 List the Nova instances:

```
[root@management-server-cisco images]# nova list
+-----+-----+-----+-----+
| ID                | Name      | Status | Task |
+-----+-----+-----+-----+
| 77ea3918-793b-4fa7-9961-10fbdc15c6e5 | cisco-vm  | ACTIVE | -    |
+-----+-----+-----+-----+
```

Step 3 Show the Nova instance ID for one of the instances:

```
[root@management-server-cisco images]# nova show
77ea3918-793b-4fa7-9961-10fbdc15c6e5 | grep instance_name
| OS-EXT-SRV-ATTR:instance_name      | instance-00000003
```

The Nova instance ID in this example is instance-00000003. This ID will be used later with the `virsh` command. Nova instance IDs are actually the libvirt IDs of the libvirt domain associated with the Nova instance.

Step 4 Identify the compute node where the VM was deployed:

```
[root@management-server-cisco images]# nova show 77ea3918-793b-4fa7-9961-10fbdc15c6e5 | grep
hypervisor
| OS-EXT-SRV-ATTR:hypervisor_hostname | compute_server-1
```

The compute node in this case is `compute_server-1`. You will connect to this compute node to call the `virsh` commands. Next, you get the volume ID from the libvirt domain in the Nova libvirt container.

Step 5 Launch a SSH session to the identified compute node, `compute_server-1`:

```
[root@management-server-cisco ~]# ssh root@compute_server-1
```

Step 6 Navigate to the Nova libvirt container:

```
[root@compute_server-1 ~]# libvirt
```

Step 7 Get the instance libvirt domain volume ID:

```
novalibvirt_4612 [root@compute_server-1 /]# virsh dumpxml instance-00000003 | grep rbd
<source protocol='rbd' name='volumes/volume-dd188a5d-f822-4769-8a57-c16694841a23'>
```

Step 8 Launch a SSH session to a controller node:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 9 Navigate to the Ceph Monitor container:

```
[root@compute_server-1 ~]# cephmon
```

Step 10 Verify volume ID matches the ID in Step 7:

```
cephmon_4612 [ceph@controller_server-1 ceph]
$ rbd list volumes | grep volume-dd188a5d-f822-4769-8a57-c16694841a23
volume-dd188a5d-f822-4769-8a57-c16694841a23
```

Displaying Docker Disk Space Usage

Docker supports multiple storage back ends such as Device Mapper, thin pool, overlay, and AUFS. Cisco VIM uses the devicemapper storage driver because it provides strong performance and thin provisioning. Device Mapper is a kernel-based framework that supports advanced volume management capability. Complete the following steps to display the disk space used by Docker containers.

Step 1 Launch a SSH session to a controller or compute node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2 Enter the docker info command to display the disk space used by Docker containers:

```
[root@controller_server_1 ~]# docker info
Containers: 24
Images: 186
Storage Driver: devicemapper
Pool Name: vg_var-docker--pool
Pool Blocksize: 524.3 kB
Backing Filesystem: xfs
Data file:
Metadata file:
Data Space Used: 17.51 GB
Data Space Total: 274.9 GB
Data Space Available: 257.4 GB...
```

Reconfiguring SwiftStack Integration

Cisco VIM provides integration with SwiftStack, an object storage solution. The key aspect of the SwiftStack integration is to add a SwiftStack endpoint to an existing pod running on Cisco VIM through the reconfigure option. In this case the SwiftStack is installed and managed outside the Cisco VIM ahead of time, and the

VIM orchestrator adds the relevant Keystone configuration details to access the SwiftStack endpoint (see the Cisco VIM install guide for more details of SwiftStack).

The following options support the SwiftStack reconfiguration:

- Enable SwiftStack integration if it is not present.
- Reconfigure the existing SwiftStack PAC endpoint to point to a different cluster (`cluster_api_endpoint`).
- Reconfigure the `Reseller_prefix` of the existing SwiftStack installation.
- Reconfigure the admin password (`admin_password`) of an existing SwiftStack Install.

Integrating SwiftStack over TLS

The automation supports SwiftStack integration over TLS. To enable TLS, the CA root certificate must be presented as part of the `/root/openstack-configs/haproxy-ca.crt` file. The protocol parameter within the SWIFTSTACK stanza must be set to `https`. As a pre-requisite, the SwiftStack cluster needs to be configured to enable HTTPS connections for the SwiftStack APIs with termination at the proxy servers.

The following section needs to be configured in the `Setup_data.yaml` file.

```
#####
# Optional Swift configuration section
#####
# SWIFTSTACK: # Identifies the objectstore provider by name
#   cluster_api_endpoint: <IP address of PAC (proxy-account-container) endpoint>
#   reseller_prefix: <Reseller_prefix as configured for Keysone Auth,AuthToken support in
Swiftstack E.g KEY_>
#   admin_user: <admin user for swift to authenticate in keystone>
#   admin_password: <swiftstack_admin_password>
#   admin_tenant: <The service tenant corresponding to the Account-Container used by
Swiftstack
#   protocol: <http or https> # protocol that swiftstack is running on top
```



Note

The operator should pay attention while updating the settings to ensure that SwiftStack over TLS are appropriately pre-configured in the customer-managed SwiftStack controller as specified in the Install guide.

To initiate the integration, copy the `setupdata` into a local directory by running the following command:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
```

Update the `setupdata` by running the following command:

```
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include SwiftStack info)
```

Run the reconfiguration command as follows:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

Cinder Volume Backup on SwiftStack

Cisco VIM enables cinder service to be configured to backup its block storage volumes to the SwiftStack object store. This feature is automatically configured if the SWIFTSTACK stanza is present in the setup_data.yaml file. The mechanism is to authenticate against SwiftStack during volume backups leverages. The same keystone SwiftStack endpoint is configured to manage objects. The default SwiftStack container that manages cinder volumes within the account (Keystone Tenant as specified by admin_tenant) is currently defaulted to volumebackups.

Reconfiguring Administrator Source Networks

To access the administrator services, Cisco VIM provides source IP based filtering of network requests on the management node. These services include SSH and Kibana dashboard access. When the services are configured all admin network requests made to the management node are dropped, except the white listed addresses in the configuration.

Reconfiguring administrator source network supports the following options:

- Set administrator source network list: Network addresses can be added or deleted from the configuration; the list is replaced in whole during a reconfigure operation.
- Remove administrator source network list: If the **admin_source_networks** option is removed, then the source address does not filter the incoming admin service requests.

The following section needs to be configured in the Setup_data.yaml file:

```
admin_source_networks: # optional, host based firewall to white list admin's source IP
- 10.0.0.0/8
- 172.16.0.0/12
```



Note

The operator has to be careful while updating the source networks. If the list is misconfigured, operators may lock themselves out of access to the management node through SSH. If it is locked, an operator must log into the management node through the console port to repair the configuration.

To initiate the integration, copy the setupdata into a local directory by running the following command:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
```

Update the setupdata by running the following command:

```
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include SwiftStack info)
```

Run the reconfiguration command as follows:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

Password Reset for Cisco VIM Management Node

Run the following command to reset the Root Password of Cisco VIM management node **RHEL-7 / systemd**

1. Boot your system and wait until the **GRUB2** menu appears.
2. In the **boot loader** menu, highlight any entry and press **e**.
3. Find the line beginning with **linux**. At the end of this line, append the following:

```
init=/bin/sh
```

Or if you face any alarm, instead of **ro** change **rw** to **sysroot** as shown in the following example:

```
rw init=/sysroot/bin/sh
```

4. Press **Ctrl+X** to boot the system using the options you edited.

Once the system boots, you can see the shell prompt without having to enter any user name or password:

```
sh-4.2#
```

5. Load the installed SELinux policy by running the following command:

```
sh-4.2# /usr/sbin/load_policy -i
```

6. Execute the following command to remount your root partition:

```
sh4.2#  
mount -o remount,rw /
```

7. Reset the root password by running the following command:

```
sh4.2# passwd root
```

When prompted, enter your new root password and click **Enter** key to confirm. Enter the password for the second time to make sure you typed it correctly and confirm with **Enter** again. If both the passwords match, a confirmation message appears.

8. Execute the following command to remount the root partition again, this time as read-only:

```
sh4.2#  
mount -o remount,ro /
```

9. Reboot the system. Now you can log in as the root user using the new password set up during this procedure.

To reboot the system, enter **exit** and **exit** again to leave the environment and reboot the system.

References: <https://access.redhat.com/solutions/918283>.



CHAPTER 6

Overview to Cisco VIM Unified Management

Cisco VIM Insight is an optional application, which acts as a single point of management for the Cisco VIM. Inclusive of your Cisco NFVI package, you can use Cisco VIM Insight to manage Cisco NFVI for day-0 and day-n and for multi-site and multi-pod management features.

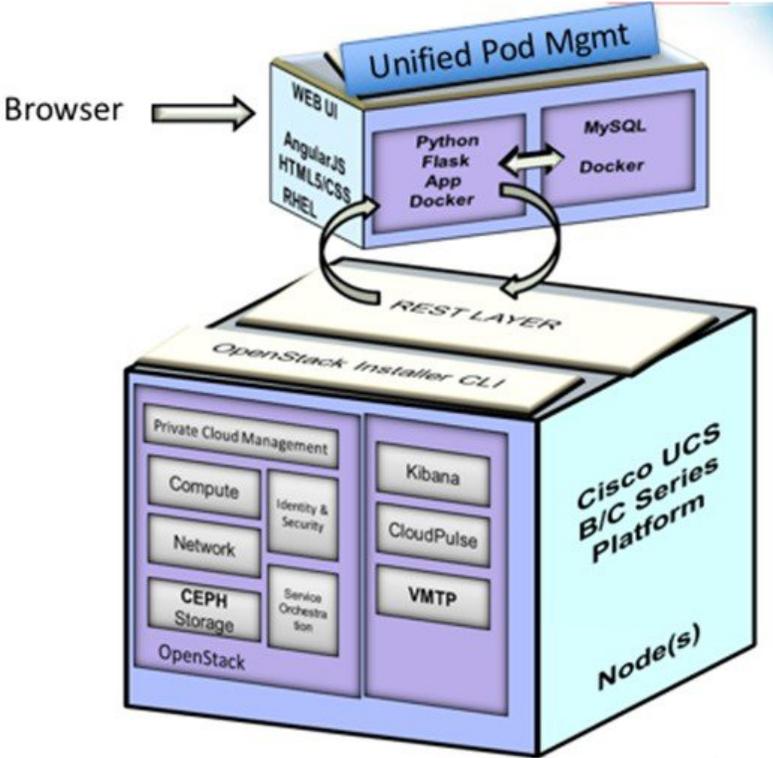
- [Cisco VIM Unified Management Overview, on page 129](#)
- [Cisco VIM Unified Management Admin UI Overview, on page 131](#)
- [Cisco VIM Unified Management Pod UI Overview, on page 131](#)

Cisco VIM Unified Management Overview

Cisco VIM provides an Intuitive and easy way to deploy and manage the NFVI platform, reducing user-error and providing visualization deployment to manage multiple Cisco VIM Pods from a single portal. In Cisco VIM 2.2 and higher releases, a light-weight UI which is a dockerized application, supports multi-tenancy with local RBAC support and CiscoVIM Rest layer are integrated. The container-based UI platform manages multiple CiscoVIM pods from day-0, or above in the lifecycle of the cloud.

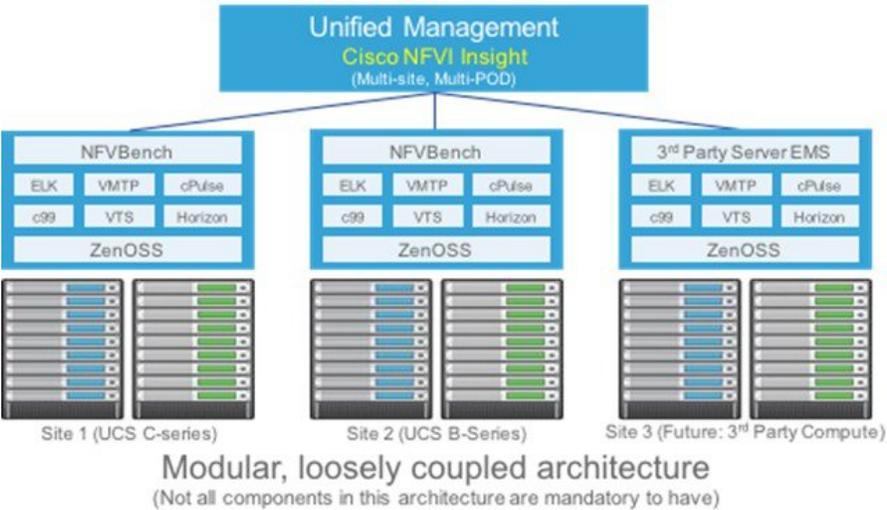
The following figure shows the architecture of the CiscoVIM UM's interaction with a Pod:

Figure 19: Cisco VIM UM's Interaction with a Pod



The architecture of the CiscoVIM UM is light-weight, hierarchical, and scalable. Each local site is autonomous with localized toolsets. Global Unified Management UI, provides ease of management with multisite and multi-pod capability for distributed NFV deployment at scale. This facility can be used through browsers such as IE, Firefox, Safari, and Chrome. Cisco VIM UM by itself, is designed to operate in HA. The platform is a modular, loosely coupled architecture, that provides the capability to manage multiple pods, with RBAC support as depicted in the following figure:

Figure 20: Cisco VIM UM Architecture



Cisco VIM UM can be installed in Standalone or non-HA mode: You can Install in a Standalone or non-HA mode (on the management node of the pod) or a standalone (BOM same as the management node) server. Migrating from one install mode to another can be done effectively as the UI interacts with each Pod through REST API and little RBAC information of the Admin and user is kept in the DB.

The UI has two types of views:

- UI Admin: UI Admin can add users as UI Admin or Pod Admin.
- Pod Admin: Pod Admin has the privilege only at the Pod level, unless Pod Admin is also a UI Admin.

Cisco VIM Unified Management Admin UI Overview

Admin UI is responsible for managing the UI and Pod admin, which includes adding and revoking user privileges. Also, the UI Admin can delete an existing Pod from the management pane.

Cisco VIM Unified Management Pod UI Overview

The Pod UI, is responsible for managing each Pod. VIM UM gives easy access to switch between multiple Pods. Through the Pod UI, a Pod Admin can manage users and their respective roles and responsibilities. Also, the Pod UI provides the user to execute day-0 (install) and day-n (Pod management, software update, and so on.) activities seamlessly. ELK, Horizon Web UI, and so on, are also cross-launched and visible for each Pod through the Pod UI.



CHAPTER 7

Managing Cisco VIM through Unified Management

This functionality brings in clear separation of roles. It does not store any pod related details obtained directly through Rest API from the pods locally, except for RBAC information.

- [UI Administrators Privileges and Responsibilities, on page 133](#)
- [Pod UI Privileges and Responsibilities, on page 134](#)
- [Adding Cisco VIM Pod, on page 134](#)
- [Deleting Pod from Cisco VIM Unified Management, on page 135](#)
- [Context Switching Within Insight, on page 136](#)

UI Administrators Privileges and Responsibilities

The Insight UI Admin has the following privileges and responsibilities:

1. Insight UI Admin(s) can only add Pod Admin.
2. Insight UI Admin can manage all the users in Insight from **Manage Pod Users**.
 - UI Admin can revoke permission of Users: If UI Admin wants to revoke a user from a Pod, click **Revoke permission** icon under Action column.
 - UI Admin can delete a User: If UI Admin wants to delete a user from the UM, Click **Delete** icon under Action column. If there is only one user associated with a Pod then UI Admin needs to delete the pod and then delete or revoke the user permission.
3. Insight UI Admin can manage Pod Admin(s) from **Manage Pod Admin**.
 - UI Admin can add a new Pod Admin in Insight.
 - UI Admin can revoke permission of a user from being a Pod Admin.
4. Insight UI Admin can manage Pods from Manage Pods.
 - UI Admin can delete a Pod from Insight.
 - UI Admin can also update password for the REST incase there was a system update on the pod and REST password was changed in that process.

5. Insight UI Admin can manage other UI Admin(s) from **Manage UI Admin Users**.
 - Insight UI Admin can add another UI Admin.
 - Insight UI Admin can revoke permission of the user from being an UI Admin.



Note If there is only one UI Admin for Insight then revoke permission icon will be disabled for the user.

Pod UI Privileges and Responsibilities

As Cisco VIM is Rest API based, you can manage a pod through CLI, Rest API or UI. You can always bring in a partial or fully functional Pod and register with VIM UM. UM queries the pod status through Rest API and reflect the same.



Note We recommended the admin to choose only one path to manage the pod.

Adding Cisco VIM Pod

Before you begin

Complete the following pre-requisites to add a Cisco VIM Pod:

- Bootstrap of VIM Insight is complete and successful as per the install guide.
- At minimum, a UI and Pod Admin exists as per the install guide.

Step 1 Navigate to https://br_api:9000

Step 2 Click **Register Management Node** link.

- Enter the Endpoint IP which is the **br_api** of your Pod.

Note Run time validation to check if the Endpoint IP is already registered to Insight.

- Give a name or tag for the pod you are registering.
- Enter the REST API password for the Pod.
 - You can locate the REST API password on the Pod you are registering.
 - The path to locate REST API password is : `/opt/cisco/ui_config.json`.
- A brief description about management node. Description field is optional and can be left blank.
- Enter the Email ID of the Pod Admin.

- Run time validation to check if the Email ID is Pod admin or not.
- If False, the Insight will give an error User is not registered as Pod Admin.
- If True, the User Name is auto-populated and the **Register** button will be enabled.

Step 3 Click **Browse** to upload restapi server CA Certificate. This is enabled once the Pod Admin validation is successful.

- Navigate to `/var/www/mercury/mercury-ca.crt` of the management node.
- Download the Certificate to the local machine and upload the certificate using Insight.

Validation check for file size and extension is done as a part of upload and in case of failure the Certificate is deleted and you need to upload the valid certificate again.

If the Certificate is uploaded successfully then **Register** button is enabled. To do a management node health check click **Register**.

- If the REST API service is down on the management node then a failure message will be displayed as : Installer REST API service not available. The certificate will not be deleted.
- If the Certificate is invalid and there is a SSL connection problem with the management node then certificate is deleted and message is displayed to upload the certificate again.
- If the Certificate is valid user is redirected to the login page with a message- management node registered successfully.

Step 4 Click **Register** to redirect the user to the landing or login page. Pod Admin receives the notification mail that the management node is registered successfully.

Deleting Pod from Cisco VIM Unified Management

When you delete a Pod from Cisco VIM UM, you are not deleting the Pod from your OpenStack deployment.

Before you begin

Following the steps to delete a Cisco VIM Pod:

- Bootstrap of VIM Insight is complete and successful as per the install guide.
- At least one UI and Pod Admin exists as per the install guide.
- The UM manages the targeted Pod.

Step 1 Log in as the **UM UI Admin**.

Step 2 In the navigation pane, click **Manage Pods**.

Step 3 Choose the pod that you want to delete in the Action column and click **Delete**.

Step 4 Click **Proceed**, to confirm the deletion.

Context Switching Within Insight

Cisco VIM UM has permissions to switch between two or more pods for a particular node. The user can be an admin for one or more pods, and a normal user for some other pod, simultaneously. Ability to access multiple pods, provides the user to maintain context and yet scale from a pod management point of view.

There are two ways a user can switch to another pod.

- **Context Switching Icon:** Context Switching Icon is situated on the top right corner and is the third icon from the right tool tip of the UI. Click **Context Switching** Icon to view all the pods that you can access. Pod with a red dot indicates that the REST Password that is entered during registration of the Management node does not match with the current REST Password for that of particular node. In such a situation the Pod admin or User has to reach out to UI admin to update the password for that Node. UI admin updates the password from Manage Pods in Insight UI admin Portal.
- **Switch Between Management Nodes:** Switch Between Management Nodes is available in the Dashboard. The user can see all the pods in the table and can navigate to any Pod using a single click. If mouse pointer changes from hand or cursor to a red dot sign it indicates that the REST Password entered during registration of Management node does not match with the current REST Password for that particular node.



CHAPTER 8

Managing Blueprints

The following topics tell you how to manage Cisco NFVI Blueprints.

- [Blueprints](#), on page 137
- [Creating a Blueprint Using Upload Functionality](#), on page 138
- [Managing Post Install Features](#), on page 176

Blueprints

Blueprints contain the configuration metadata required to deploy an OpenStack system through a Cisco VIM pod in Cisco VIM Unified Management. You can create a blueprint in Cisco UM or you can upload a yaml file that contains the metadata for a blueprint. You can also create a blueprint from an existing OpenStack system that you are configuring as a Cisco VIM pod.

The configuration in the blueprint is specific to the type of Cisco UCS server that is in the OpenStack system. A blueprint for a C-Series server-based OpenStack system cannot be used to configure a B-Series server-based OpenStack system. Cisco UM displays an error if the blueprint does not match the configuration of the OpenStack system.

The blueprint enables you to quickly change the configuration of an OpenStack system. While only one blueprint can be active, you can create or upload multiple blueprints for a Cisco VIM pod. If you change the active blueprint for a pod, you have to update the configuration of the OpenStack system to match the new blueprint.



Note You can modify and validate an existing blueprint, or delete a blueprint. However, you cannot modify any of the configuration metadata in the active blueprint for a Cisco VIM pod.

Blueprint Activation

A blueprint becomes active when you use it in a successful installation for a Cisco VIM pod. Other blueprints that you created or uploaded to that pod are in nonactive state.

Uploading or creating a blueprint does not activate that blueprint for the pod. Install a blueprint through the **Cisco VIM Suite** wizard. If the installation is successful, the selected blueprint becomes active.



Note If you want to activate a new blueprint in an existing pod, you have to delete certain accounts and the credential policies for that pod before you activate the blueprint. See [Activating a Blueprint in an Existing Pod with OpenStack Installed, on page 139](#).

Viewing Blueprint Details

To view blueprint details:

-
- Step 1** Log in to Cisco VIM Insight as pod user.
 - Step 2** Choose the Cisco VIM pod with the blueprint that you want to view.
 - Step 3** Click **Menu** at the top left corner to expand the navigation pane.
 - Step 4** Choose **Pre-Install > Blueprint Management**.
 - Step 5** Choose a blueprint from the list.
 - Step 6** Click **Preview and Download YAML**.
-

Creating a Blueprint Using Upload Functionality

Before you begin

- You must have a YAML file (B series or C Series) on your system.
- Only one blueprint can be uploaded at a time. To create a blueprint off-line, refer to the `setup_data.yaml.B_Series_EXAMPLE` or `setup_data.yaml.C_Series_EXAMPLE`.
- The respective keys in the sample YAML have to match or the corresponding pane does not get populated during the upload.

-
- Step 1** Log in to **Cisco VIM UM**.
 - Step 2** In the navigation pane, expand the **Pre-Install** section and click **Blueprint** setup.
 - Step 3** Click the **Browse** in the **Blueprint Initial Setup**.
 - Step 4** Click **Select**.
 - Step 5** Click **Load** in the **Insight UI Application**.
All the fields present in the YAML file is uploaded to the respective fields in the UI.
 - Step 6** Provide a **Name for the Blueprint**.
While saving the blueprint name has to be unique.
 - Step 7** Click **Offline Validation**.

- If all the mandatory fields in the UI are populated, then Offline Validation of the Blueprint commences, or else a pop up message indicating the section of Blueprint creation that has missing information error shows up.

Step 8 On Offline Blueprint Validation being successful, **Save Blueprint** and **Cancel** is enabled.

Note If the Blueprint Validation Fails, only the **Cancel** button is enabled.

Activating a Blueprint in an Existing Pod with OpenStack Installed

Before you begin

You must have a POD which has an active Installation of OpenStack. If the OpenStack installation is in Failed State, then UM UI will not be able to fetch the Blueprint.

Step 1 Go to the **Landing page** of the UM Log in.

Step 2 Click **Register Management Node**.

Step 3 Enter the following details:

- Management Node IP Address.
- Management Node Name (Any friendly Name).
- REST API Password (/opt/cisco/ui_config.json).
- Description about the Management Node.
- POD Admin's Email ID.

A notification email is sent to the email id entered during registration.

Step 4 Log in using the same email id and password.

Step 5 In the navigation pane, click **Pre-Install > Blueprint Management**.

Choose the **NEWSETUPDATA** from the **Blueprint Management** pane.

This is the same setup data which was used by ciscovimclient, to run the installation on the Management Node.

Blueprint Management



Note You must have at least one blueprint (In any state Active or In-Active or In-progress), in the Blueprint Management Pane.

The screenshot shows the Cisco VIM Unified Management interface. The top navigation bar includes the Cisco logo, 'VIM Unified Management', the user's name 'Calsoft', IP address '10.30.116.244', role 'Role: Full-Pod-Access', and user ID 'User ID: rohassar'. A notification bar indicates 'Press F11 to exit full screen'. The left sidebar contains navigation options: Dashboard, Pre-Install, Blueprint Setup (with sub-option 'Blueprint Management' selected), Post-Install, View Topology, and Pod User Administration. The main content area is titled 'Blueprints Management' and features a table with columns for 'Blueprint Title', 'Modified Date', 'Status', and 'Action'. The table lists four blueprints: 'Test' (Invalid), '5555' (Invalid), 'NEWSETUPDATA' (Deployed), and '56646' (Invalid). Each row has an 'Action' column with icons for edit, delete, and download. A search bar and pagination controls are also visible.

Blueprint Title	Modified Date	Status	Action
Test	4/3/2018, 2:55:18 PM	Invalid	[Edit] [Delete] [Download]
5555	4/2/2018, 9:27:07 PM	Invalid	[Edit] [Delete] [Download]
NEWSETUPDATA	4/3/2018, 5:15:25 PM	Deployed	[Edit] [Delete] [Download]
56646	4/2/2018, 9:29:00 PM	Invalid	[Edit] [Delete] [Download]

Blueprint Management grid contains the list of all the blueprints that are saved. You can save the blueprint even if it is failed in the Blueprint Setup. However, you will not be allowed to deploy those Blueprints.

Blueprint Management table provides the following information:

- Blueprint Name
- Modified Date
- Edit, Remove, and Download Blueprint
- Search Blueprint

Blueprint Name: It shows the name of the Blueprint. You cannot edit this field. It shows the name of the blueprint that is saved after Offline Validation.



Note No two blueprints can have the same Blueprint name.

Modified Date: This shows when blueprint was last modified.

Blueprint Status: There are 6 total status for the Blueprint.

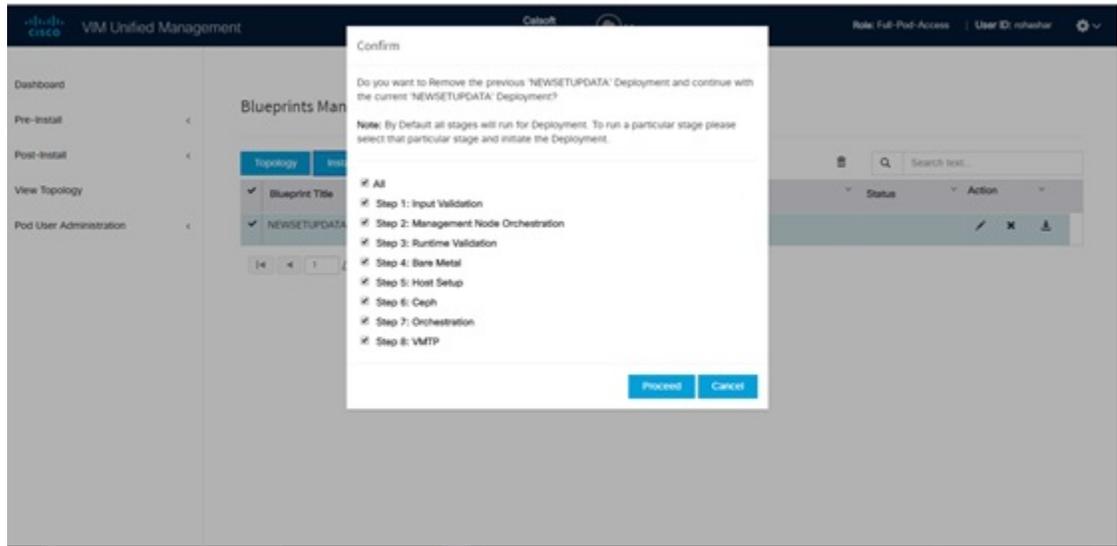
- Valid: Blueprint that is saved after offline validation success.
- Invalid: Blueprint that is saved after Offline Validation failure.
- Inprogress: Blueprint that is saved without running Offline Validation.
- Deployed: Blueprint that is used to bring up cloud without failures.
- Installing: Blueprint that is used to initiate the cloud deployment.

- Failed: Blueprint that is used to deploy the cloud which eventually failed.

With every blueprint record, there are some operations associated that you can perform by using the buttons – Topology, Install, and Remove.

Topology

Topology allows you to view graphical representation of the control, compute, and storage node that is associated with the various network segments.



Install Button

Click **Install**, a confirmation message is generated requesting to initiate the deployment with the stages you want to run. By default all stages are selected but you can also do an incremented install. In case of Incremented Install, you have to choose stages in the order. For Example: If you choose Validation Stage then the 2nd stage Management Node Orchestration is enabled. You cannot skip stages and run a deployment. Once you click **Proceed**, the Cloud Deployment is initiated and the progress can be viewed from the Dashboard.

Remove Button

Choose the blueprint and click **Remove** to remove the blueprint. A confirmation message appears. If you click **Proceed**, the blueprint removal operation is initiated.

Edit, Remove, and Download Blueprint

You can edit or delete a Blueprint which is not in Deployed State. If you want to take a backup of the Blueprint locally, click *Download* icon which generates the preview to download the Blueprint.

Following are the ways to deploy a Blueprint:

- If there is no Blueprint in Deployed state, then you can choose any Valid Blueprint from the list.
- If there is a Blueprint in a Failed state, then you can choose another Valid Blueprint but Insight asks you to remove the previous deployment before proceeding.
- If there is a Blueprint in Deployed state, then you can choose another Valid Blueprint but Insight asks you to remove the previous deployment before proceeding.

The deployment of Blueprint occurs stepwise and if any one step fails for some reason, a **Play** button is displayed on that particular step. You can click a **Play** button and begin the installation for that particular state.



Note There is always one blueprint in Deployed state. You cannot deploy multiple blueprints in the cloud.

Search Blueprint: Search box is displayed on top-right of the table which facilitates you to lookup for Blueprint by their name or status. Navigate to **Topology** and choose a Blueprint which redirects you to the default blueprint, the one which is selected in the Blueprint Management pane.



Note During the various operations across the application the cloud icon in the center of the header changes its color which is based on the following table.

Table 3:

POD Operation	Status	Icon or Color
Management Node Registered, No Active Deployment	Pending	Gray
Cloud Up And Running, No Failure	Active	Green
Cloud Installation/ Any Operation In Progress	In-Progress	Blue
Cloudpulse Failed	Critical Warnings	Red
Pod Operation Failed	Warning	Amber
Software Update (Auto) Rollback Failed	Critical Warnings	Red
Uncommitted Software Update	Warning	Amber
Reconfigure Openstack Password	Critical Warning	Red
Reconfigure CIMC Password	Warning	Amber
Reconfigure Optional Features/ OS	Critical Warning	Red
Power Management Operation Fails	Warning	Amber
Management Not-Reachable	Not-Reachable	Red

Creating a Blueprint for B-Series Server Platform

Typically, you create the blueprint when you create the Cisco VIM pod. Follow the instructions below to create an additional blueprint for a pod that uses B-Series servers.

Before you begin

Create a Cisco VIM Insight User Account and Register the respective Pod.

Step 1

Log-in to Cisco VIM Insight.

Step 2

In the **Navigation** pane, expand the **Pre-Install Section**.

Step 3

Click **Blueprint Setup**.

Step 4

On the **Blueprint Initial Setup** page of the Cisco VIM Insight, complete the following fields:

Name	Description
Blueprint Name field	Enter the name for the blueprint configuration.
Platform Type drop-down list	Choose one of the following platform types: <ul style="list-style-type: none"> • B-Series (By Default) • C-Series
Tenant Network drop-down list	Choose one of the following tenant network types: <ul style="list-style-type: none"> • Linux Bridge/VXLAN • OVS/VLAN
Ceph Mode drop-down list	Choose one of the following Ceph types: <ul style="list-style-type: none"> • Dedicated • Central (By Default) (not supported in production)
Pod Type drop-down list	Fullon (By default).
Optional Features and Services checkbox	Syslog Export Settings, Swiftstack, Nfvbench, VMTP, LDAP, Pod Name, TOR Switch Information, TLS, Heat, Vim Admins, Auto Backup, NFVI Monitoring, Install Mode, Keystone v3, Enable Esc Priv. If any one is selected, the corresponding section is visible in various Blueprint sections. By default all options are disabled.
Import Existing YAML file field	If you have an existing B Series YAML file you can use this feature to upload the file. Insight will automatically fill in the fields and if any mandatory fields are missed then the respective section will be highlighted.

Step 5

Click **Physical Setup** to advance to the **Registry Setup** configuration page. Fill in the following details for Registry Setup:

Name	Description
Registry User Name text field	User-Name for Registry (Mandatory).
Registry Password text field	Password for Registry (Mandatory).
Registry Email text field	Email ID for Registry (Mandatory).

Once all mandatory fields are filled the **Validation Check Registry** page will be changed to a Green Tick.

Step 6

Click **UCSM Common** tab and fill the following fields:

Name	Description
User name disabled field	By default value is admin.
Password text field	Enter Password for UCSM Common (Mandatory).
UCSM IP text field	Enter IP Address for UCSM Common (Mandatory).
Resource Prefix text field	Enter the resource prefix (Mandatory)
QOS Policy Type drop-down list	Choose one of the following types: <ul style="list-style-type: none"> • NFVI (Default) • Media
Enable Prov FI PIN optional checkbox	Default is false.
MRAID-CARD optional checkbox	Enables JBOD mode to be set on disks. Applicable only if you have RAID controller configured on Storage C240 Rack servers.
Enable UCSM Plugin optional checkbox	Visible when Tenant Network type is OVS/VLA.
Enable QoS Policy optional checkbox	Visible only when UCSM Plugin is enabled. If UCSM Plugin is disabled then this option will be set to False.
SRIOV Multi VLAN Trunk optional grid	Visible when UCSM Plugin is enabled. Enter the values for network and vlans ranges. Grid can handle all CRUD operations like Add, Delete, Edit and Multiple Delete.

Step 7

Click **Networking** to advance to the networking section of the Blueprint.

Name	Description
Domain Name field	Enter the domain name (Mandatory).
HTTP Proxy Server field	If your configuration uses an HTTP proxy server, enter the IP address of the server.
HTTPS Proxy Server field	If your configuration uses an HTTPS proxy server, enter the IP address of the server.

Name	Description
IP Tables on Management Pods	
NTP Servers field	Enter a maximum of four and minimum of one IPv4 and/or IPv6 addresses in the table.
Domain Name Servers field	Enter a maximum of three and minimum of one IPv4 and/or IPv6 addresses.

Name	Description
Network table	

Name	Description												
	<p>Network table is pre-populated with segments. To add Networks you can either clear all the table using Delete all or click Edit icon for each segment and fill in the details.</p> <p>You can add, edit, or delete network information in the table.</p> <ul style="list-style-type: none"> • Click Add to enter new entries (networks) to the table. • Specify the following fields in the Edit Entry to Networks dialog: <table border="1" data-bbox="911 653 1531 1837"> <tr> <td data-bbox="911 653 1219 808">VALN field</td> <td data-bbox="1219 653 1531 808">Enter the VLAN ID. For Segment - Provider, the VLAN ID value is always none.</td> </tr> <tr> <td data-bbox="911 808 1219 1430">Segment drop-down list</td> <td data-bbox="1219 808 1531 1430"> You can select any of one segment from dropdown list <ul style="list-style-type: none"> • API • Management Provision • Tenant • CIMC • Storage • External • Provider (optional) <p>Note Depending upon the segment not all entries listed below are needed</p> </td> </tr> <tr> <td data-bbox="911 1430 1219 1518">Subnet field</td> <td data-bbox="1219 1430 1531 1518">Enter the IPv4 address for the subnet.</td> </tr> <tr> <td data-bbox="911 1518 1219 1673">IPv6 Subnet field</td> <td data-bbox="1219 1518 1531 1673">Enter IPv6 Subnet Address. This field will be available only for Management Provision and API .</td> </tr> <tr> <td data-bbox="911 1673 1219 1766">Gateway field</td> <td data-bbox="1219 1673 1531 1766">Enter the IPv4 address for the Gateway.</td> </tr> <tr> <td data-bbox="911 1766 1219 1837">IPv6 Gateway field</td> <td data-bbox="1219 1766 1531 1837">Enter IPv6 gateway. This field is only available for the</td> </tr> </table>	VALN field	Enter the VLAN ID. For Segment - Provider, the VLAN ID value is always none .	Segment drop-down list	You can select any of one segment from dropdown list <ul style="list-style-type: none"> • API • Management Provision • Tenant • CIMC • Storage • External • Provider (optional) <p>Note Depending upon the segment not all entries listed below are needed</p>	Subnet field	Enter the IPv4 address for the subnet.	IPv6 Subnet field	Enter IPv6 Subnet Address. This field will be available only for Management Provision and API .	Gateway field	Enter the IPv4 address for the Gateway.	IPv6 Gateway field	Enter IPv6 gateway. This field is only available for the
VALN field	Enter the VLAN ID. For Segment - Provider, the VLAN ID value is always none .												
Segment drop-down list	You can select any of one segment from dropdown list <ul style="list-style-type: none"> • API • Management Provision • Tenant • CIMC • Storage • External • Provider (optional) <p>Note Depending upon the segment not all entries listed below are needed</p>												
Subnet field	Enter the IPv4 address for the subnet.												
IPv6 Subnet field	Enter IPv6 Subnet Address. This field will be available only for Management Provision and API .												
Gateway field	Enter the IPv4 address for the Gateway.												
IPv6 Gateway field	Enter IPv6 gateway. This field is only available for the												

Name	Description	
		Management Provision and API .
	Pool field	Pool can be defined with single IP, range of IP or discontinuous pool. Enter the pool information in the required format: Single IP: Example: 10.30.118.101 Range of IP: Example: 10.30.118.98 to 10.30.118.105 Discontinuous IP: Example: 10.30.118.101, 10.30.118.98 to 10.30.118.105
	IPv6 Pool field	Enter the pool information in the required format. For Example: 10.1.15-10.1.1.10,10.2.15-10.2.1.10 This field is available only for Management Provision.
Click Save .		

Step 8

On the **Servers and Roles** page of the **Cisco VIM Suite** wizard, click **Add (+)** to add a new entry in the table, and complete the following fields:

Name	Description
Server User Name	Enter the username of the Server.
Disable Hyperthreading	Default value is false. You can set it as true or false.

Name	Description																	
Cobbler	Enter the Cobbler details in the following fields:																	
	<table border="1"> <thead> <tr> <th data-bbox="909 344 1213 386">Name</th> <th data-bbox="1219 344 1523 386">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="909 386 1219 596"> Cobbler Timeout field </td> <td data-bbox="1219 386 1523 596"> The default value is 45 min. This is an optional parameter. Timeout is displayed in minutes, and its value ranges from 30 to 120. </td> </tr> <tr> <td data-bbox="909 596 1219 684"> Block Storage Kickstart field </td> <td data-bbox="1219 596 1523 684"> Kickstart file for Storage Node. </td> </tr> <tr> <td data-bbox="909 684 1219 894"> Admin Password Hash field </td> <td data-bbox="1219 684 1523 894"> Enter the Admin Password. Password should be Alphanumeric. Password should contain minimum 8 characters and maximum of 32 characters. </td> </tr> <tr> <td data-bbox="909 894 1219 982"> Cobbler Username field </td> <td data-bbox="1219 894 1523 982"> Enter the cobbler username to access the cobbler server. </td> </tr> <tr> <td data-bbox="909 982 1219 1071"> Control Kickstart field </td> <td data-bbox="1219 982 1523 1071"> Kickstart file for Control Node. </td> </tr> <tr> <td data-bbox="909 1071 1219 1159"> Compute Kickstart field </td> <td data-bbox="1219 1071 1523 1159"> Kickstart file for Compute Node. </td> </tr> <tr> <td data-bbox="909 1159 1219 1247"> Cobbler Admin Username field </td> <td data-bbox="1219 1159 1523 1247"> Enter the admin username of the Cobbler. </td> </tr> </tbody> </table>	Name	Description	Cobbler Timeout field	The default value is 45 min. This is an optional parameter. Timeout is displayed in minutes, and its value ranges from 30 to 120.	Block Storage Kickstart field	Kickstart file for Storage Node.	Admin Password Hash field	Enter the Admin Password. Password should be Alphanumeric. Password should contain minimum 8 characters and maximum of 32 characters.	Cobbler Username field	Enter the cobbler username to access the cobbler server.	Control Kickstart field	Kickstart file for Control Node.	Compute Kickstart field	Kickstart file for Compute Node.	Cobbler Admin Username field	Enter the admin username of the Cobbler.	
	Name	Description																
	Cobbler Timeout field	The default value is 45 min. This is an optional parameter. Timeout is displayed in minutes, and its value ranges from 30 to 120.																
	Block Storage Kickstart field	Kickstart file for Storage Node.																
	Admin Password Hash field	Enter the Admin Password. Password should be Alphanumeric. Password should contain minimum 8 characters and maximum of 32 characters.																
	Cobbler Username field	Enter the cobbler username to access the cobbler server.																
	Control Kickstart field	Kickstart file for Control Node.																
Compute Kickstart field	Kickstart file for Compute Node.																	
Cobbler Admin Username field	Enter the admin username of the Cobbler.																	

Name	Description	
<p>Add Entry to Servers and Roles.</p>	<p>Click Edit or + to add a new server and role to the table.</p>	
	<p>Server Name</p>	<p>Enter a server name.</p>
	<p>Server Type drop-down list.</p>	<p>Choose Blade or Rack from the drop-down list.</p>
	<p>Rack ID field.</p>	<p>The Rack ID for the server.</p>
	<p>Chassis ID field</p>	<p>Enter a Chassis ID.</p>
	<p>If Rack is chosen, the Rack Unit ID field is displayed.</p>	<p>Enter a Rack Unit ID.</p>
	<p>If Blade is chosen, the Blade ID field is displayed.</p>	<p>Enter a Blade ID.</p>
	<p>Select the Role from the drop down list.</p>	<p>If Server type is Blade then Control and Compute. If Rack is selected then Block Storage.</p>
	<p>Management IP field.</p>	<p>It is an optional field but if provided for one server then it is mandatory to provide it for other Servers as well.</p>
<p>Management IPv6 field.</p>	<p>Enter Management Ipv6 address.</p>	
<p>Click Save or Add.</p>	<p>Clicking Save or Add, adds all information for Servers and Roles.</p>	

Step 9

Click **ToR Switch** checkbox in Blueprint Initial Setup to enable the **TOR SWITCH** configuration page. It is an **Optional** section in Blueprint Setup, but once all the fields are filled in then it will become a part of the Blueprint.

Name	Description
<p>Configure ToR optional checkbox .</p>	<p>If you enable this checkbox, the Configure ToR section will change from false to true.</p>

Name	Description	
<p>ToR Switch Information mandatory table if you want to enter ToR information.</p>	Click + to add information for ToR Switch.	
	Name	Description
	Name	ToR switch name.
	Username	ToR switch username.
	Password	ToR switch Password.
	SSH IP	ToR switch SSH IP Address.
	SSN Num	ToR switch ssn num. output of show license host-id.
	VPC Peer Keepalive	Peer Management IP. You need not define if there is no peer as it is optional but it will become mandatory when the ToR is in VPC.
	VPC Domain	Need not define if there is no peer.
	VPC Peer port	Interface for vpc peer ports.
	VPC Peer VLAN Info	vlan ids for vpc peer ports (optional).
BR Management Port Info	Management interface of build node.	
BR Management PO Info	Port channel number for management interface of build node.	
On clicking Save , Add ToR Info connected to Fabric field will be visible.	Port Channel field.	Enter the port channel input.
	Switch Name field.	Enter the switch name.

Step 10 Click **OpenStack Setup** tab to advance to the **OpenStack Setup** Configuration page.

Step 11 On the **OpenStack Setup** page of the Cisco VIM Insight wizard, complete the following fields:

Name	Description	
HA Proxy	Fill in the mandatory fields:	
	External VIP Address	Enter IP address of External VIP.
	External VIP Address IPv6	Enter IPv6 address of External VIP.
	Virtual Router ID	Enter the Router ID for HA.
	Internal VIP Address IPv6	Enter IPv6 address.
	Internal VIP Address	Enter IP address of Internal VIP.
Keystone	Mandatory fields are pre-populated. This option is always true.	
	Admin Username	admin.
	Admin Tenant Name	admin.

Name	Description																												
<p>LDAP on Keystone.</p> <p>Note: this option is only available with Keystone v3</p>	<p>This is available only when Keystone v3 and LDAP both are enabled under Optional Features and Services in Blueprint Initial Setup.</p> <table border="1" data-bbox="867 405 1529 1392"> <tr> <td data-bbox="867 405 1198 459">Domain Name field</td> <td data-bbox="1198 405 1529 459">Enter name for Domain name.</td> </tr> <tr> <td data-bbox="867 459 1198 514">Object Class for Users field</td> <td data-bbox="1198 459 1529 514">Enter a string as input.</td> </tr> <tr> <td data-bbox="867 514 1198 569">Object Class for Groups</td> <td data-bbox="1198 514 1529 569">Enter a string.</td> </tr> <tr> <td data-bbox="867 569 1198 661">Domain Name Tree for Users</td> <td data-bbox="1198 569 1529 661">Enter a string.</td> </tr> <tr> <td data-bbox="867 661 1198 753">Domain Name Tree for Groups field</td> <td data-bbox="1198 661 1529 753">Enter a string.</td> </tr> <tr> <td data-bbox="867 753 1198 846">Suffix for Domain Name field</td> <td data-bbox="1198 753 1529 846">Enter a string.</td> </tr> <tr> <td data-bbox="867 846 1198 938">URL field</td> <td data-bbox="1198 846 1529 938">Enter a URL with ending port number.</td> </tr> <tr> <td data-bbox="867 938 1198 1031">Domain Name for Bind User field</td> <td data-bbox="1198 938 1529 1031">Enter a string.</td> </tr> <tr> <td data-bbox="867 1031 1198 1123">Password field</td> <td data-bbox="1198 1031 1529 1123">Enter Password as string format.</td> </tr> <tr> <td data-bbox="867 1123 1198 1169">User Filter field</td> <td data-bbox="1198 1123 1529 1169">Enter filter name as string.</td> </tr> <tr> <td data-bbox="867 1169 1198 1224">User ID Attribute field</td> <td data-bbox="1198 1169 1529 1224">Enter a string.</td> </tr> <tr> <td data-bbox="867 1224 1198 1278">User Name Attribute field</td> <td data-bbox="1198 1224 1529 1278">Enter a string.</td> </tr> <tr> <td data-bbox="867 1278 1198 1333">User Mail Attributefield</td> <td data-bbox="1198 1278 1529 1333">Enter a string.</td> </tr> <tr> <td data-bbox="867 1333 1198 1388">Group Name Attribute field</td> <td data-bbox="1198 1333 1529 1388">Enter a string.</td> </tr> </table>	Domain Name field	Enter name for Domain name.	Object Class for Users field	Enter a string as input.	Object Class for Groups	Enter a string.	Domain Name Tree for Users	Enter a string.	Domain Name Tree for Groups field	Enter a string.	Suffix for Domain Name field	Enter a string.	URL field	Enter a URL with ending port number.	Domain Name for Bind User field	Enter a string.	Password field	Enter Password as string format.	User Filter field	Enter filter name as string.	User ID Attribute field	Enter a string.	User Name Attribute field	Enter a string.	User Mail Attributefield	Enter a string.	Group Name Attribute field	Enter a string.
Domain Name field	Enter name for Domain name.																												
Object Class for Users field	Enter a string as input.																												
Object Class for Groups	Enter a string.																												
Domain Name Tree for Users	Enter a string.																												
Domain Name Tree for Groups field	Enter a string.																												
Suffix for Domain Name field	Enter a string.																												
URL field	Enter a URL with ending port number.																												
Domain Name for Bind User field	Enter a string.																												
Password field	Enter Password as string format.																												
User Filter field	Enter filter name as string.																												
User ID Attribute field	Enter a string.																												
User Name Attribute field	Enter a string.																												
User Mail Attributefield	Enter a string.																												
Group Name Attribute field	Enter a string.																												

Name	Description														
Neutron	<p>Neutron fields change on the basis of Tenant Network Type Selection from Blueprint Initial Setup page.</p> <p>Following are the options available for Neutron for OVS/VLAN:</p> <table border="1" data-bbox="829 451 1485 1486"> <tr> <td data-bbox="829 451 1159 604">Tenant Network Type</td> <td data-bbox="1159 451 1485 604">Auto Filled based on the Tenant Network Type selected in the Blueprint Initial Setup page.</td> </tr> <tr> <td data-bbox="829 604 1159 724">Mechanism Drivers</td> <td data-bbox="1159 604 1485 724">Auto Filled based on the Tenant Network Type selected in Blueprint Initial Setup page.</td> </tr> <tr> <td data-bbox="829 724 1159 1129">NFV Hosts</td> <td data-bbox="1159 724 1485 1129"> Auto filled with the Compute you added in Server and Roles. If you select All in this section NFV_HOSTS: ALL will be added to the Blueprint or you can select one particular compute. For Eg: NFV_HOSTS: compute-server-1, compute-server-2. </td> </tr> <tr> <td data-bbox="829 1129 1159 1220">Tenant VLAN Ranges</td> <td data-bbox="1159 1129 1485 1220">List of ranges separated by comma of form start:end.</td> </tr> <tr> <td data-bbox="829 1220 1159 1310">Provider VLAN Ranges</td> <td data-bbox="1159 1220 1485 1310">List of ranges separated by comma of form start:end.</td> </tr> <tr> <td data-bbox="829 1310 1159 1430">VM Hugh Page Size (available for NFV_HOSTS option)</td> <td data-bbox="1159 1310 1485 1430">2M or 1G</td> </tr> <tr> <td data-bbox="829 1430 1159 1486">Enable Jumbo Frames</td> <td data-bbox="1159 1430 1485 1486">Check Box</td> </tr> </table> <p>For Tenant Network Type Linux Bridge, everything will remain the same except Tenant VLAN Ranges which will be removed.</p>	Tenant Network Type	Auto Filled based on the Tenant Network Type selected in the Blueprint Initial Setup page.	Mechanism Drivers	Auto Filled based on the Tenant Network Type selected in Blueprint Initial Setup page.	NFV Hosts	Auto filled with the Compute you added in Server and Roles. If you select All in this section NFV_HOSTS: ALL will be added to the Blueprint or you can select one particular compute. For Eg: NFV_HOSTS: compute-server-1, compute-server-2.	Tenant VLAN Ranges	List of ranges separated by comma of form start:end.	Provider VLAN Ranges	List of ranges separated by comma of form start:end.	VM Hugh Page Size (available for NFV_HOSTS option)	2M or 1G	Enable Jumbo Frames	Check Box
Tenant Network Type	Auto Filled based on the Tenant Network Type selected in the Blueprint Initial Setup page.														
Mechanism Drivers	Auto Filled based on the Tenant Network Type selected in Blueprint Initial Setup page.														
NFV Hosts	Auto filled with the Compute you added in Server and Roles. If you select All in this section NFV_HOSTS: ALL will be added to the Blueprint or you can select one particular compute. For Eg: NFV_HOSTS: compute-server-1, compute-server-2.														
Tenant VLAN Ranges	List of ranges separated by comma of form start:end.														
Provider VLAN Ranges	List of ranges separated by comma of form start:end.														
VM Hugh Page Size (available for NFV_HOSTS option)	2M or 1G														
Enable Jumbo Frames	Check Box														
CEPH	<p>Ceph has two pre-populated fields</p> <ul style="list-style-type: none"> • CEPH Mode: By default Dedicated. • NOVA Boot from: From the drop-down, choose Ceph or local. 														

Name	Description																										
GLANCE	By default Populated for CEPH Dedicated with Store Backend value as CEPH .																										
CINDER	By default Populated for CEPH Dedicated with Volume Driver value as CEPH .																										
<p>VMTP optional section will only be visible once VMTP is selected from Blueprint Initial Setup.</p>	<p>Check one of the check boxes to specify a VMTP network:</p> <ul style="list-style-type: none"> • Provider Network • External Network <p>For the Provider Network complete the following:</p> <table border="1" data-bbox="867 682 1528 1276"> <tr> <td data-bbox="867 682 1195 768">Network Name field.</td> <td data-bbox="1195 682 1528 768">Enter the name for the external network.</td> </tr> <tr> <td data-bbox="867 768 1195 854">IP Start field.</td> <td data-bbox="1195 768 1528 854">Enter the starting floating IPv4 address.</td> </tr> <tr> <td data-bbox="867 854 1195 940">IP End field.</td> <td data-bbox="1195 854 1528 940">Enter the ending floating IPv4 address.</td> </tr> <tr> <td data-bbox="867 940 1195 1026">Gateway field</td> <td data-bbox="1195 940 1528 1026">Enter the IPv4 address for the Gateway.</td> </tr> <tr> <td data-bbox="867 1026 1195 1113">DNS Server field.</td> <td data-bbox="1195 1026 1528 1113">Enter the DNS server IPv4 address.</td> </tr> <tr> <td data-bbox="867 1113 1195 1178">Segmentation ID field.</td> <td data-bbox="1195 1113 1528 1178">Enter the segmentation ID.</td> </tr> <tr> <td data-bbox="867 1178 1195 1276">Subnet</td> <td data-bbox="1195 1178 1528 1276">Enter the Subnet for Provider Network.</td> </tr> </table> <p>For External Network fill in the following details:</p> <table border="1" data-bbox="867 1346 1528 1862"> <tr> <td data-bbox="867 1346 1195 1432">Network Name field.</td> <td data-bbox="1195 1346 1528 1432">Enter the name for the external network.</td> </tr> <tr> <td data-bbox="867 1432 1195 1518">Network IP Start field.</td> <td data-bbox="1195 1432 1528 1518">Enter the starting floating IPv4 address.</td> </tr> <tr> <td data-bbox="867 1518 1195 1604">Network IP End field.</td> <td data-bbox="1195 1518 1528 1604">Enter the ending floating IPv4 address.</td> </tr> <tr> <td data-bbox="867 1604 1195 1690">Network Gateway field</td> <td data-bbox="1195 1604 1528 1690">Enter the IPv4 address for the Gateway.</td> </tr> <tr> <td data-bbox="867 1690 1195 1776">DNS Server field.</td> <td data-bbox="1195 1690 1528 1776">Enter the DNS server IPv4 address.</td> </tr> <tr> <td data-bbox="867 1776 1195 1862">Subnet</td> <td data-bbox="1195 1776 1528 1862">Enter the Subnet for External Network.</td> </tr> </table>	Network Name field.	Enter the name for the external network.	IP Start field.	Enter the starting floating IPv4 address.	IP End field.	Enter the ending floating IPv4 address.	Gateway field	Enter the IPv4 address for the Gateway.	DNS Server field.	Enter the DNS server IPv4 address.	Segmentation ID field.	Enter the segmentation ID.	Subnet	Enter the Subnet for Provider Network.	Network Name field.	Enter the name for the external network.	Network IP Start field.	Enter the starting floating IPv4 address.	Network IP End field.	Enter the ending floating IPv4 address.	Network Gateway field	Enter the IPv4 address for the Gateway.	DNS Server field.	Enter the DNS server IPv4 address.	Subnet	Enter the Subnet for External Network.
Network Name field.	Enter the name for the external network.																										
IP Start field.	Enter the starting floating IPv4 address.																										
IP End field.	Enter the ending floating IPv4 address.																										
Gateway field	Enter the IPv4 address for the Gateway.																										
DNS Server field.	Enter the DNS server IPv4 address.																										
Segmentation ID field.	Enter the segmentation ID.																										
Subnet	Enter the Subnet for Provider Network.																										
Network Name field.	Enter the name for the external network.																										
Network IP Start field.	Enter the starting floating IPv4 address.																										
Network IP End field.	Enter the ending floating IPv4 address.																										
Network Gateway field	Enter the IPv4 address for the Gateway.																										
DNS Server field.	Enter the DNS server IPv4 address.																										
Subnet	Enter the Subnet for External Network.																										

Name	Description												
<p>TLS section will be visible if TLS is selected from Blueprint Initial Setup Page.</p>	<p>TLS has two options:</p> <ul style="list-style-type: none"> • External LB VIP FQDN - Text Field. • External LB VIP TLS - True/False. By default this option is false. 												
<p>SwiftStack optional section will be visible if SwiftStack is selected from Blueprint Initial Setup Page. SwiftStack is only supported with KeyStonev2 . If you select Keystonev3, swiftstack cannot be configured.</p>	<p>Following are the options that needs to be filled for SwiftStack:</p> <table border="1" data-bbox="828 541 1484 1171"> <tbody> <tr> <td data-bbox="828 541 1159 663">Cluster End Point</td> <td data-bbox="1159 541 1484 663">IP address of PAC (proxy-account-container) endpoint.</td> </tr> <tr> <td data-bbox="828 663 1159 751">Admin User</td> <td data-bbox="1159 663 1484 751">Admin user for swift to authenticate in keystone.</td> </tr> <tr> <td data-bbox="828 751 1159 907">Admin Tenant</td> <td data-bbox="1159 751 1484 907">The service tenant corresponding to the Account-Container used by Swiftstack.</td> </tr> <tr> <td data-bbox="828 907 1159 1062">Reseller Prefix</td> <td data-bbox="1159 907 1484 1062">Reseller_prefix as configured for Keysone Auth,AuthToken support in Swiftstack E.g KEY_</td> </tr> <tr> <td data-bbox="828 1062 1159 1117">Admin Password</td> <td data-bbox="1159 1062 1484 1117">swiftstack_admin_password</td> </tr> <tr> <td data-bbox="828 1117 1159 1171">Protocol</td> <td data-bbox="1159 1117 1484 1171">http or https</td> </tr> </tbody> </table>	Cluster End Point	IP address of PAC (proxy-account-container) endpoint.	Admin User	Admin user for swift to authenticate in keystone.	Admin Tenant	The service tenant corresponding to the Account-Container used by Swiftstack.	Reseller Prefix	Reseller_prefix as configured for Keysone Auth,AuthToken support in Swiftstack E.g KEY_	Admin Password	swiftstack_admin_password	Protocol	http or https
Cluster End Point	IP address of PAC (proxy-account-container) endpoint.												
Admin User	Admin user for swift to authenticate in keystone.												
Admin Tenant	The service tenant corresponding to the Account-Container used by Swiftstack.												
Reseller Prefix	Reseller_prefix as configured for Keysone Auth,AuthToken support in Swiftstack E.g KEY_												
Admin Password	swiftstack_admin_password												
Protocol	http or https												
<p>Under the openstack setup tab, the Vim_admins tab will only be visible once Vim_admins is selected from the Optional Features & Services under the Blueprint InitialSetup tab.</p>	<p>Following are the options that needs to be filled for Vim Admins:</p> <ul style="list-style-type: none"> • Username - Text Field • Password - Password field. Admin hash password should always start with \$6 												

Step 12

If **Syslog Export** or **NFVBENCH** is selected in **Blueprint Initial Setup** Page, then **Services Setup** page would be **enabled** for user to view. Following are the options under **Services Setup Tab**:

Name	Description										
Syslog Export.	Following are the options for Syslog Settings:										
	<table border="1"> <tr> <td>Remote Host</td> <td>Enter Syslog IP Address</td> </tr> <tr> <td>Facility</td> <td>Defaults to local5</td> </tr> <tr> <td>Severity</td> <td>Defaults to debug</td> </tr> <tr> <td>Clients</td> <td>Defaults to ELK</td> </tr> <tr> <td>Port</td> <td>Defaults to 514 but can be modified by the User.</td> </tr> </table>	Remote Host	Enter Syslog IP Address	Facility	Defaults to local5	Severity	Defaults to debug	Clients	Defaults to ELK	Port	Defaults to 514 but can be modified by the User.
	Remote Host	Enter Syslog IP Address									
	Facility	Defaults to local5									
	Severity	Defaults to debug									
	Clients	Defaults to ELK									
Port	Defaults to 514 but can be modified by the User.										
NFVBENCH	<p>Enable checkbox which by default is False.</p> <p>Add Tor information connected to switch:</p> <ul style="list-style-type: none"> • Select a TOR Switch and Enter the Switch name. • Enter the port number. For example: eth1/5. VTEP VLANS (mandatory and needed only for VXLAN): Enter 2 different VLANs for VLAN1 and VLAN2. • NIC Ports: INT1 and INT2 optional input. Enter the 2 port numbers of the 4-port 10G Intel NIC at the management node used for NFVBench. 										

Step 13 Click **Offline validation** to initiate an offline Blueprint validation.

Step 14 Once the **Offline validation** is successful, **Save** option will be enabled which will redirect you to the **Blueprint Management** page.

Creating a Blueprint for C-Series Server Platform

Create a Cisco VIM Insight User Account and register the respective Pod.

Step 1 Log-in to **CISCO VIM Insight**.

Step 2 In the **Navigation** pane, expand the **Pre-Install Section**.

Step 3 Click **Blueprint Setup**.

Step 4 On the **Blueprint Initial Setup** page of the Cisco VIM Insight , complete the following fields:

Name	Description
Blueprint Name field	Enter the name for the blueprint configuration.
Platform Type drop-down list	<ul style="list-style-type: none"> • B-Series (By Default) • C-Series (Select C Series)

Name	Description
<p>Tenant Network drop-down list</p>	<p>Choose one of the following tenant network types:</p> <ul style="list-style-type: none"> • Linux Bridge/VXLAN • OVS/VLAN • VTS/VLAN • VPP/VLAN • ACI/VLAN <p>Note when VTS/VLAN or ACI/VLAN is selected then respective tabs are available on Blueprint setup</p>
<p>Pod Type drop-down list</p>	<p>Choose one of the following pod type :</p> <ul style="list-style-type: none"> • Fullon(By Default) • Micro • UMHC <p>Note UMHC pod type is only supported for OVS/VLAN tenant type.</p> <p>Note Pod type micro is supported for OVS/VLAN, ACI/VLAN,VPP/VLAN.</p>
<p>Ceph Mode drop-down list</p>	<p>Choose one of the following Ceph types:</p> <ul style="list-style-type: none"> • Dedicated (By Default) • Central (Is not supported in production)
<p>Optional Features and Services checkbox.</p>	<p>Swiftstack, LDAP, Syslog Export Settings, Install Mode, TorSwitch Information, TLS, Nfvmon, Pod Name, VMTP, Nfvbench, Auto Backup, Heat, Keystone v3, Enable Esc Priv.</p> <p>If any one is selected, the corresponding section is visible in various Blueprint sections.</p> <p>By default all options are disabled.</p>
<p>Import Existing YAML file</p>	<p>If you have an existing C Series YAML file you can use this feature to upload the file.</p> <p>Insight will automatically fill in the fields and if any mandatory field is missed then would highlight it in the respective section.</p>

Step 5 Click **Physical Setup** to advance to the **Registry Setup** configuration page. Fill in the following details for Registry Setup.

Name	Description
Registry User Name text field	User-Name for Registry (Mandatory).
Registry Password text field	Password for Registry (Mandatory).
Registry Email text field	Email ID for Registry (Mandatory).

Once all Mandatory fields are filled, the **Validation Check Registry** page will indicate a green tick.

Step 6

Click **CIMC Common** tab and complete the following fields:

Name	Description
User Name disabled field	By default value is Admin.
Password text field	Enter Password for UCSM Common (Mandatory).

Step 7

Click **Networking** to advance to the networking section of the Blueprint.

Name	Description
Domain Name field.	Enter the domain name (Mandatory) .
NTP Servers field.	Enter a maximum of four and minimum of one IPv4 and/or IPv6 addresses in the table.
Domain Name Servers field	Enter a maximum of three and minimum of one IPv4 and/or IPv6 addresses
HTTP Proxy Server field	If your configuration uses an HTTP proxy server, enter the IP address of the server.
HTTPS Proxy Server field.	If your configuration uses an HTTPS proxy server, enter the IP address of the server.

Name	Description
Networks table	

Name	Description														
	<p>Network table is pre-populated with segments. To add Networks you can either clear all the table using Delete all or click Edit icon for each segment and fill in the details.</p> <p>You can add, edit, or delete network information in the table.</p>														
	<table border="1"> <thead> <tr> <th data-bbox="906 506 1219 554">Name</th> <th data-bbox="1219 506 1534 554">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="906 554 1219 722"> VLAN field </td> <td data-bbox="1219 554 1534 722"> Enter the VLAN ID. For Segment - Provider, the VLAN ID value is always "none". </td> </tr> <tr> <td data-bbox="906 722 1219 1409"> Segment drop-down list </td> <td data-bbox="1219 722 1534 1409"> You can select any one segment from the dropdown list. <ul style="list-style-type: none"> • API • Management/Provision • Tenant • CIMC • Storage • External • Provider (optional) <p>Note Some segments do not need some of the values listed in the preceding points.</p> </td> </tr> <tr> <td data-bbox="906 1409 1219 1499"> Subnet field </td> <td data-bbox="1219 1409 1534 1499"> Enter the IPv4 address for the subnet. </td> </tr> <tr> <td data-bbox="906 1499 1219 1654"> IPv6 Subnet field </td> <td data-bbox="1219 1499 1534 1654"> Enter IPv6 address. This field will be available only for Management provision and API. </td> </tr> <tr> <td data-bbox="906 1654 1219 1738"> Gateway field </td> <td data-bbox="1219 1654 1534 1738"> Enter the IPv4 address for the Gateway. </td> </tr> <tr> <td data-bbox="906 1738 1219 1841"> IPv6 Gateway field </td> <td data-bbox="1219 1738 1534 1841"> Enter IPv6 gateway. This field will only available only for Management </td> </tr> </tbody> </table>	Name	Description	VLAN field	Enter the VLAN ID. For Segment - Provider, the VLAN ID value is always "none".	Segment drop-down list	You can select any one segment from the dropdown list. <ul style="list-style-type: none"> • API • Management/Provision • Tenant • CIMC • Storage • External • Provider (optional) <p>Note Some segments do not need some of the values listed in the preceding points.</p>	Subnet field	Enter the IPv4 address for the subnet.	IPv6 Subnet field	Enter IPv6 address. This field will be available only for Management provision and API.	Gateway field	Enter the IPv4 address for the Gateway.	IPv6 Gateway field	Enter IPv6 gateway. This field will only available only for Management
Name	Description														
VLAN field	Enter the VLAN ID. For Segment - Provider, the VLAN ID value is always "none".														
Segment drop-down list	You can select any one segment from the dropdown list. <ul style="list-style-type: none"> • API • Management/Provision • Tenant • CIMC • Storage • External • Provider (optional) <p>Note Some segments do not need some of the values listed in the preceding points.</p>														
Subnet field	Enter the IPv4 address for the subnet.														
IPv6 Subnet field	Enter IPv6 address. This field will be available only for Management provision and API.														
Gateway field	Enter the IPv4 address for the Gateway.														
IPv6 Gateway field	Enter IPv6 gateway. This field will only available only for Management														

Name	Description									
	<table border="1"> <thead> <tr> <th data-bbox="870 287 1177 331">Name</th> <th data-bbox="1177 287 1484 331">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="870 331 1177 382"></td> <td data-bbox="1177 331 1484 382">provision and API network.</td> </tr> <tr> <td data-bbox="870 382 1177 533"> Pool field </td> <td data-bbox="1177 382 1484 533">Enter the pool information in the required format, for example: 10.30.1.1 or 10.30.1.1 to 10.30.1.12</td> </tr> <tr> <td data-bbox="870 533 1177 766"> IPv6 Pool field </td> <td data-bbox="1177 533 1484 766">Enter the pool information in the required format, for example: 10.1.1.5-10.1.1.10,10.2.1.5-10.2.1.10 This field is only available for the Mgmt/Provision.</td> </tr> </tbody> </table>	Name	Description		provision and API network.	Pool field	Enter the pool information in the required format, for example: 10.30.1.1 or 10.30.1.1 to 10.30.1.12	IPv6 Pool field	Enter the pool information in the required format, for example: 10.1.1.5-10.1.1.10,10.2.1.5-10.2.1.10 This field is only available for the Mgmt/Provision.	
	Name	Description								
		provision and API network.								
Pool field	Enter the pool information in the required format, for example: 10.30.1.1 or 10.30.1.1 to 10.30.1.12									
IPv6 Pool field	Enter the pool information in the required format, for example: 10.1.1.5-10.1.1.10,10.2.1.5-10.2.1.10 This field is only available for the Mgmt/Provision.									
	provision and API network.									
Pool field	Enter the pool information in the required format, for example: 10.30.1.1 or 10.30.1.1 to 10.30.1.12									
IPv6 Pool field	Enter the pool information in the required format, for example: 10.1.1.5-10.1.1.10,10.2.1.5-10.2.1.10 This field is only available for the Mgmt/Provision.									

Step 8

On the **Servers and Roles** page of the **Cisco VIM Suite** wizard, click **Add (+)** to add a new entry in the table, and complete the following fields:

You can edit or delete existing entries in the **Server and Roles** table.

Name	Description
Server User Name	Enter the username of the Server.
Disable Hyperthreading	Default value is false. You can set it as true or false.

Name	Description	
Cobbler	Enter the Cobbler details in the following fields:	
	Name	Description
	Cobbler Timeout field	The default value is 45 min. This is an optional parameter. Timeout is displayed in minutes, and its value ranges from 30 to 120.
	Block Storage Kickstart field	Kickstart file for Storage Node.
	Admin Password Hash field	Enter the Admin Password. Password should be Alphanumeric. Password should contain minimum 8 characters and maximum of 32 characters.
	Cobbler Username field	Enter the cobbler username to access the cobbler server.
	Control Kickstart field	Kickstart file for Control Node.
	Compute Kickstart field	Kickstart file for Compute Node.
Cobbler Admin Username field	Enter the admin username of the Cobbler.	

Name	Description																			
<p>Add Entry to Servers and Roles .</p> <p>Note when Pod type micro is selected then all the three servers will be associated with control, compute and block storage role.</p> <p>For Example:</p> <p>Roles</p> <ul style="list-style-type: none"> • Block Storage <ul style="list-style-type: none"> • -Server 1 • -Server 2 • -Server 3 • Control <ul style="list-style-type: none"> • -Server 1 • -Server 2 • -Server 3 • Compute <ul style="list-style-type: none"> • -Server 1 • -Server 2 • -Server 3 <p>Note When Pod type UMHC is selected then auto ToR configuration is not supported and the ToR info at server and roles level is not allowed to be entered.</p>	<p>Click Edit or + to add a new server and role to the table.</p> <table border="1" data-bbox="776 338 1487 1045"> <tr> <td data-bbox="776 338 1133 394">Server Name</td> <td data-bbox="1133 338 1487 394">Enter the server name .</td> </tr> <tr> <td data-bbox="776 394 1133 451">Rack ID field</td> <td data-bbox="1133 394 1487 451">The rack ID for the server.</td> </tr> <tr> <td data-bbox="776 451 1133 508">VIC Slot field</td> <td data-bbox="1133 451 1487 508">Enter a VIC Slot.</td> </tr> <tr> <td data-bbox="776 508 1133 598">Management IPv6field</td> <td data-bbox="1133 508 1487 598">This is optional field. Enter Ipv6 format address</td> </tr> <tr> <td data-bbox="776 598 1133 655">CIMC IP field</td> <td data-bbox="1133 598 1487 655">Enter a IP address.</td> </tr> <tr> <td data-bbox="776 655 1133 711">CIMC Username field</td> <td data-bbox="1133 655 1487 711">Enter a Username.</td> </tr> <tr> <td data-bbox="776 711 1133 768">CIMC Password field</td> <td data-bbox="1133 711 1487 768">Enter a Password for CIMC</td> </tr> <tr> <td data-bbox="776 768 1133 890">Select the Role from the drop down list</td> <td data-bbox="1133 768 1487 890">Choose Control or Compute or Block Storage from the drop-down list.</td> </tr> <tr> <td data-bbox="776 890 1133 1045">Management IP</td> <td data-bbox="1133 890 1487 1045">It is an optional field but if provided for one server then it is mandatory to provide it for other servers.</td> </tr> </table>		Server Name	Enter the server name .	Rack ID field	The rack ID for the server.	VIC Slot field	Enter a VIC Slot.	Management IPv6 field	This is optional field. Enter Ipv6 format address	CIMC IP field	Enter a IP address.	CIMC Username field	Enter a Username.	CIMC Password field	Enter a Password for CIMC	Select the Role from the drop down list	Choose Control or Compute or Block Storage from the drop-down list.	Management IP	It is an optional field but if provided for one server then it is mandatory to provide it for other servers.
Server Name	Enter the server name .																			
Rack ID field	The rack ID for the server.																			
VIC Slot field	Enter a VIC Slot.																			
Management IPv6 field	This is optional field. Enter Ipv6 format address																			
CIMC IP field	Enter a IP address.																			
CIMC Username field	Enter a Username.																			
CIMC Password field	Enter a Password for CIMC																			
Select the Role from the drop down list	Choose Control or Compute or Block Storage from the drop-down list.																			
Management IP	It is an optional field but if provided for one server then it is mandatory to provide it for other servers.																			
Click Save or Add .	On clicking Save or Add all information related to Servers and Roles gets saved.																			
If Configure ToR checkbox is True with at-least one switch detail, these fields will be displayed for each server and this is similar to DP Tor: Port Channel and Switch Name (Mandatory if Configure ToR is true)	<ul style="list-style-type: none"> • Port Channel field • Switch Name field • Switch Port Info field 	<ul style="list-style-type: none"> • Enter the port channel input. • Enter the switch name. • Enter the switch port information. 																		

Name	Description	
DP ToR (Only for Control and Compute) : Mandatory if Intel NIC and Configure TOR is True.	<ul style="list-style-type: none"> • Port Channel field • Switch Name field • Switch Port Info field 	<ul style="list-style-type: none"> • Enter the port channel input. • Enter the switch name. • Enter the switch port information.
SRIOV TOR INFO (Only for Compute Nodes). It is mandatory in server and roles if Intel NIC and Configure TOR is True. Switch Name (Mandatory if Configure ToR is true) . This field appears only when Intel NIC support is true, as Auto TOR config is not supported in VIC_NIC combo	<ul style="list-style-type: none"> • Switch Name field • Switch Port Info field 	<ul style="list-style-type: none"> • Enter the switch name. • Enter the switch port information.
Intel SRIOV VFS (valid for Intel NIC testbeds) and can be integer.	For SRIOV support for Intel NIC. By Default, SRIOV support is disabled. To enable, define a value in the range # * 1-32 when INTEL_NIC_SUPPORT is set True (X710 Max VFs = 32) # * 1-63 when CISCO_VIC_INTEL_SRIOV is set True (X520 Max VFs = 63)	
INTEL_SRIOV_PHYS_PORTS (valid for Intel NIC test beds) and can be of value 2 or 4 (default is 2)	In some cases the # of Physical SRIOV port needed is 4; to meet that requirement, define the following: # this is optional, if nothing is defined code will assume it to be 2; the only 2 integer values this parameter # takes is 2 or 4 and is true when INTEL_NIC_SUPPORT is True and INTEL_SRIOV_VFS is valid	
Click Save or Add .	On clicking Save or Add all information related to Servers and Roles gets saved.	

Note Maximum two ToR info needs to be configured for each connection type on each node (control, compute and block_storage node).

Note If pod type UMHC is selected then CISCO_VIC_INTEL_SRIOV is enabled to be TRUE.

Note For Tenant type ACI/VLAN, port channel for each ToR port will not be available in servers and roles, as APIC will automatically assign port-channel numbers.

Step 9 Click **ToR Switch** checkbox in **Blueprint Initial Setup** to enable the **TOR SWITCH** configuration page. It is an **Optional** section in Blueprint Setup but once all the fields are filled, it becomes a part of the Blueprint.

Name	Description	
<p>Configure TOR optional checkbox.</p> <p>Note If UMHC is selected as podtype, configure TOR is not allowed.</p>	<p>If you enable this checkbox configure tor section would be changed from false to true.</p> <p>Note Configure tor is true then ToR switch info maps in servers</p>	

Name	Description																								
<p>TOR Switch Information mandatory table if you want to enter ToR information.</p>	Click + to add information for ToR Switch.																								
	<table border="1"> <thead> <tr> <th data-bbox="777 342 1133 394">Name</th> <th data-bbox="1133 342 1479 394">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="777 394 1133 453">Name</td> <td data-bbox="1133 394 1479 453">ToR Switch Name.</td> </tr> <tr> <td data-bbox="777 453 1133 512">Username</td> <td data-bbox="1133 453 1479 512">TOR switch username</td> </tr> <tr> <td data-bbox="777 512 1133 571">Password</td> <td data-bbox="1133 512 1479 571">ToR switch Password</td> </tr> <tr> <td data-bbox="777 571 1133 630">SSH IP</td> <td data-bbox="1133 571 1479 630">TOR switch ssh ip</td> </tr> <tr> <td data-bbox="777 630 1133 688">SSN Num</td> <td data-bbox="1133 630 1479 688">TOR switch ssn num</td> </tr> <tr> <td data-bbox="777 688 1133 772">VPC Peer Keepalive</td> <td data-bbox="1133 688 1479 772">Peer Management IP. Do not define if there is no peer</td> </tr> <tr> <td data-bbox="777 772 1133 831">VPC Domain</td> <td data-bbox="1133 772 1479 831">Do not define if there is no peer</td> </tr> <tr> <td data-bbox="777 831 1133 890">VPC Peer Port Info</td> <td data-bbox="1133 831 1479 890">Interface for vpc peer ports</td> </tr> <tr> <td data-bbox="777 890 1133 974">VPC Peer VLAN Info</td> <td data-bbox="1133 890 1479 974">vlan ids for vpc peer ports (optional)</td> </tr> <tr> <td data-bbox="777 974 1133 1058">BR Management Port Info</td> <td data-bbox="1133 974 1479 1058">Management interface of build node</td> </tr> <tr> <td data-bbox="777 1058 1133 1192">BR Management PO Info</td> <td data-bbox="1133 1058 1479 1192">Port channel number for management interface of build node</td> </tr> </tbody> </table>	Name	Description	Name	ToR Switch Name.	Username	TOR switch username	Password	ToR switch Password	SSH IP	TOR switch ssh ip	SSN Num	TOR switch ssn num	VPC Peer Keepalive	Peer Management IP. Do not define if there is no peer	VPC Domain	Do not define if there is no peer	VPC Peer Port Info	Interface for vpc peer ports	VPC Peer VLAN Info	vlan ids for vpc peer ports (optional)	BR Management Port Info	Management interface of build node	BR Management PO Info	Port channel number for management interface of build node
	Name	Description																							
	Name	ToR Switch Name.																							
	Username	TOR switch username																							
	Password	ToR switch Password																							
	SSH IP	TOR switch ssh ip																							
	SSN Num	TOR switch ssn num																							
	VPC Peer Keepalive	Peer Management IP. Do not define if there is no peer																							
	VPC Domain	Do not define if there is no peer																							
	VPC Peer Port Info	Interface for vpc peer ports																							
	VPC Peer VLAN Info	vlan ids for vpc peer ports (optional)																							
BR Management Port Info	Management interface of build node																								
BR Management PO Info	Port channel number for management interface of build node																								
Click Save .																									

Note When tenant type ACI/VLAN is selected, the TOR switch information table differs and is mandatory.

Name	Description
Configure ToR	Is not checked, as by default ACI will configure the ToRs

	Click + to add information for ToR Switch	
	Name	Description
	Host Name	ToR switch name.
	VPC Peerkeep alive	Peer info must exist in pair.
	VPC Domain	Enter an Integer.
	BR Management Port Info	Enter Br management port info eg. Eth1/19, must have a pair in the peer switch.
Enter Node ID	Entered Integer must be unique.	

Note If TOR_TYPE is selected as NCS-5500, the TOR switch information table differs and is mandatory

Name	Description
<p>Configure ToR optional checkbox.</p> <p>Note If NSC-5500 is selected as TOR_TYPE, configure TOR is set as mandatory.</p>	<p>Enabling this checkbox, changes the configure ToR section from false to true.</p> <p>Note Configure TOR is true then ToR switchinfo maps in servers.</p>

Name	Description																										
If you want to enter Fretta details fill in the NCS-5500 Information table.	Click (+) to add information for Fretta Switch.																										
	<table border="1"> <thead> <tr> <th data-bbox="872 344 1182 392">Name</th> <th data-bbox="1182 344 1474 392">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="872 392 1182 485">Name</td> <td data-bbox="1182 392 1474 485">Enter the NCS-5500 hostname.</td> </tr> <tr> <td data-bbox="872 485 1182 577">User Name</td> <td data-bbox="1182 485 1474 577">Enter the NCS-5500 username.</td> </tr> <tr> <td data-bbox="872 577 1182 669">Password</td> <td data-bbox="1182 577 1474 669">Enter the NCS-5500 password.</td> </tr> <tr> <td data-bbox="872 669 1182 762">SSH IP</td> <td data-bbox="1182 669 1474 762">Enter the NCS-5500 ssh IP Address.</td> </tr> <tr> <td data-bbox="872 762 1182 854">VPC Peer Link</td> <td data-bbox="1182 762 1474 854">Peer management IP.</td> </tr> <tr> <td data-bbox="872 854 1182 947">BR Management PO Info</td> <td data-bbox="1182 854 1474 947">Port channel number for management interface of build node.</td> </tr> <tr> <td data-bbox="872 947 1182 1039">BR Management VLAN info</td> <td data-bbox="1182 947 1474 1039">VLAN ID for management interface of build node (access).</td> </tr> <tr> <td data-bbox="872 1039 1182 1131">VPC Peer Port Info</td> <td data-bbox="1182 1039 1474 1131">Interface for vpc peer ports.</td> </tr> <tr> <td data-bbox="872 1131 1182 1224">VPC Peer Port Address</td> <td data-bbox="1182 1131 1474 1224">Address for ISIS exchange.</td> </tr> <tr> <td data-bbox="872 1224 1182 1316">ISIS Loopback Interface address</td> <td data-bbox="1182 1224 1474 1316">ISIS loopback IP Address.</td> </tr> <tr> <td data-bbox="872 1316 1182 1388">ISIS net entity title</td> <td data-bbox="1182 1316 1474 1388">Enter a String.</td> </tr> <tr> <td data-bbox="872 1388 1182 1394">ISIS prefix SID</td> <td data-bbox="1182 1388 1474 1394">Integer between 16000 to 1048575.</td> </tr> </tbody> </table>	Name	Description	Name	Enter the NCS-5500 hostname.	User Name	Enter the NCS-5500 username.	Password	Enter the NCS-5500 password.	SSH IP	Enter the NCS-5500 ssh IP Address.	VPC Peer Link	Peer management IP.	BR Management PO Info	Port channel number for management interface of build node.	BR Management VLAN info	VLAN ID for management interface of build node (access).	VPC Peer Port Info	Interface for vpc peer ports.	VPC Peer Port Address	Address for ISIS exchange.	ISIS Loopback Interface address	ISIS loopback IP Address.	ISIS net entity title	Enter a String.	ISIS prefix SID	Integer between 16000 to 1048575.
	Name	Description																									
	Name	Enter the NCS-5500 hostname.																									
	User Name	Enter the NCS-5500 username.																									
	Password	Enter the NCS-5500 password.																									
	SSH IP	Enter the NCS-5500 ssh IP Address.																									
	VPC Peer Link	Peer management IP.																									
	BR Management PO Info	Port channel number for management interface of build node.																									
	BR Management VLAN info	VLAN ID for management interface of build node (access).																									
	VPC Peer Port Info	Interface for vpc peer ports.																									
	VPC Peer Port Address	Address for ISIS exchange.																									
	ISIS Loopback Interface address	ISIS loopback IP Address.																									
ISIS net entity title	Enter a String.																										
ISIS prefix SID	Integer between 16000 to 1048575.																										

When ToR-TYPE selected as NCS-5500 and 2 NCS-5500 are configured it is mandatory to configure MULTI_SEGMENT_ROUTING_INFO.

Name	Description
BGP AS Number	Integer between 1 to 65535.
ISIS Area Tag	A valid string.
Loopback Interface name	Loopback Interface name.
API bundle ID	Integer between 1 to 65535.

Name	Description
API bridge domain	String (Optional, only needed when br_api of mgmt node is also going through NCS-5500; this item and api_bundle_id are mutually exclusive).
EXT bridge domain	A valid string (user pre-provisions physical, bundle interface, sub-interface and external BD for external uplink and provides external BD info setup_data).

Step 10 Click **OpenStack Setup** Tab to advance to the **OpenStack Setup** page.

Step 11 In the **OpenStack Setup** page of the Cisco VIM Insight wizard, complete the following fields:

Name	Description												
Neutron	Neutron fields would change on the basis of Tenant Network Type Selection from Blueprint Initial Setup . Following are the options available for Neutron:												
	<table border="1"> <tr> <td>Tenant Network Type</td> <td>Auto Filled based on the Tenant Network Type selection in Blueprint Initial Setup page.</td> </tr> <tr> <td>Mechanism Drivers</td> <td>Auto Filled based on the Tenant Network Type selection in Blueprint Initial Setup page.</td> </tr> <tr> <td>NFV Hosts</td> <td>Auto filled with the Compute you added in Server and Roles. If you select All in this section NFV_HOSTS: "ALL" will be added to the Blueprint or else you can select particular computes as well for eg: NFV_HOSTS: "compute-server-1, compute-server-2"</td> </tr> <tr> <td>Tenant VLAN Ranges</td> <td>Allowed with VTS/VLAN VPP/VLAN, OVS/VLAN, ACI/VLAN</td> </tr> <tr> <td>Enable Jumbo Frames</td> <td>Check Box default is false.</td> </tr> <tr> <td>Huge page size Note : . This is available only when Compute node is present in NFV host</td> <td>The following are the drop-downs: <ul style="list-style-type: none"> • 2M • 1G </td> </tr> </table>	Tenant Network Type	Auto Filled based on the Tenant Network Type selection in Blueprint Initial Setup page.	Mechanism Drivers	Auto Filled based on the Tenant Network Type selection in Blueprint Initial Setup page.	NFV Hosts	Auto filled with the Compute you added in Server and Roles. If you select All in this section NFV_HOSTS: "ALL" will be added to the Blueprint or else you can select particular computes as well for eg: NFV_HOSTS: "compute-server-1, compute-server-2"	Tenant VLAN Ranges	Allowed with VTS/VLAN VPP/VLAN, OVS/VLAN, ACI/VLAN	Enable Jumbo Frames	Check Box default is false.	Huge page size Note : . This is available only when Compute node is present in NFV host	The following are the drop-downs: <ul style="list-style-type: none"> • 2M • 1G
	Tenant Network Type	Auto Filled based on the Tenant Network Type selection in Blueprint Initial Setup page.											
	Mechanism Drivers	Auto Filled based on the Tenant Network Type selection in Blueprint Initial Setup page.											
	NFV Hosts	Auto filled with the Compute you added in Server and Roles. If you select All in this section NFV_HOSTS: "ALL" will be added to the Blueprint or else you can select particular computes as well for eg: NFV_HOSTS: "compute-server-1, compute-server-2"											
	Tenant VLAN Ranges	Allowed with VTS/VLAN VPP/VLAN, OVS/VLAN, ACI/VLAN											
	Enable Jumbo Frames	Check Box default is false.											
Huge page size Note : . This is available only when Compute node is present in NFV host	The following are the drop-downs: <ul style="list-style-type: none"> • 2M • 1G 												
For Tenant Network Type Linux Bridge everything will remain the same but Tenant VLAN Ranges will be removed.													

Name	Description											
CEPH	Ceph has two pre-populated fields <ul style="list-style-type: none"> • CEPH Mode : By default Dedicated. • NOVA Boot from: Drop Down selection. You can choose Ceph or local. 											
GLANCE	By default populated for CEPH Dedicated with Store Backend value as CEPH .											
CINDER	By default Populated for CEPH Dedicated with Volume Driver value as CEPH .											
HA Proxy	Enter the Mandatory fields: <table border="1" data-bbox="834 709 1484 1115"> <tbody> <tr> <td data-bbox="834 709 1159 793">External VIP Address</td> <td data-bbox="1164 709 1484 793">Enter IP Address of External VIP.</td> </tr> <tr> <td data-bbox="834 800 1159 884">External VIP Address IPv6</td> <td data-bbox="1164 800 1484 884">Enter IP v6 Address of External VIP .</td> </tr> <tr> <td data-bbox="834 890 1159 942">Virtual Router ID</td> <td data-bbox="1164 890 1484 942">Enter the Router ID for HA.</td> </tr> <tr> <td data-bbox="834 949 1159 1033">Internal VIP Address</td> <td data-bbox="1164 949 1484 1033">Enter IP Address of Internal VIP.</td> </tr> <tr> <td data-bbox="834 1039 1159 1123">Internal VIP Address IPv6</td> <td data-bbox="1164 1039 1484 1123">Enter IP v6 Address for Internal VIP.</td> </tr> </tbody> </table>		External VIP Address	Enter IP Address of External VIP.	External VIP Address IPv6	Enter IP v6 Address of External VIP .	Virtual Router ID	Enter the Router ID for HA.	Internal VIP Address	Enter IP Address of Internal VIP.	Internal VIP Address IPv6	Enter IP v6 Address for Internal VIP.
External VIP Address	Enter IP Address of External VIP.											
External VIP Address IPv6	Enter IP v6 Address of External VIP .											
Virtual Router ID	Enter the Router ID for HA.											
Internal VIP Address	Enter IP Address of Internal VIP.											
Internal VIP Address IPv6	Enter IP v6 Address for Internal VIP.											
Keystone	<table border="1" data-bbox="834 1142 1484 1241"> <tbody> <tr> <td data-bbox="834 1142 1159 1194">Admin Username</td> <td data-bbox="1164 1142 1484 1194">admin</td> </tr> <tr> <td data-bbox="834 1201 1159 1253">Admin Tenant Name</td> <td data-bbox="1164 1201 1484 1253">admin</td> </tr> </tbody> </table>		Admin Username	admin	Admin Tenant Name	admin						
Admin Username	admin											
Admin Tenant Name	admin											

Name	Description	
<p>LDAP</p>	<p>This is available only when Keystone v3 and LDAP both are enabled under Optional Features and Services in Blueprint Initial Setup.</p>	
	<p>Domain Name field</p>	<p>Enter name for Domain name.</p>
	<p>Object Class for Users field</p>	<p>Enter a string as input.</p>
	<p>Object Class for Groups</p>	<p>Enter a string.</p>
	<p>Domain Name Tree for Users</p>	<p>Enter a string.</p>
	<p>Domain Name Tree for Groups field</p>	<p>Enter a string.</p>
	<p>Suffix for Domain Name field</p>	<p>Enter a string.</p>
	<p>URL field</p>	<p>Enter a URL with ending port number.</p>
	<p>Domain Name for Bind User field</p>	<p>Enter a string.</p>
	<p>Password field</p>	<p>Enter Password as string format.</p>
	<p>User Filter</p>	<p>Enter filter name as string.</p>
	<p>User ID Attribute</p>	<p>Enter a string.</p>
	<p>User Name Attribute</p>	<p>Enter a string.</p>
	<p>User Mail Attribute</p>	<p>Enter a string.</p>
<p>Group Name Attribute</p>	<p>Enter a string.</p>	

Name	Description																										
<p>VMTP optional section will only be visible once VMTP is selected from Blueprint Initial Setup.</p> <p>Note For VTS, Provider network is only supported</p>	<p>Check one of the check boxes to specify a VMTP network:</p> <ul style="list-style-type: none"> • Provider Network • External Network <p>For the Provider Network complete the following:</p> <table border="1" data-bbox="834 506 1484 1094"> <tbody> <tr> <td data-bbox="834 506 1159 590">Network Name field</td> <td data-bbox="1164 506 1484 590">Enter the name for the external network.</td> </tr> <tr> <td data-bbox="834 596 1159 680">IP Start field</td> <td data-bbox="1164 596 1484 680">Enter the starting floating IPv4 address.</td> </tr> <tr> <td data-bbox="834 686 1159 770">IP End field</td> <td data-bbox="1164 686 1484 770">Enter the ending floating IPv4 address.</td> </tr> <tr> <td data-bbox="834 777 1159 861">Gateway field</td> <td data-bbox="1164 777 1484 861">Enter the IPv4 address for the Gateway.</td> </tr> <tr> <td data-bbox="834 867 1159 951">DNS Server field</td> <td data-bbox="1164 867 1484 951">Enter the DNS server IPv4 address.</td> </tr> <tr> <td data-bbox="834 957 1159 1041">Segmentation ID field</td> <td data-bbox="1164 957 1484 1041">Enter the segmentation ID.</td> </tr> <tr> <td data-bbox="834 1047 1159 1094">Subnet</td> <td data-bbox="1164 1047 1484 1094">Enter the Subnet for Provider Network.</td> </tr> </tbody> </table> <p>For External Network fill in the following details:</p> <table border="1" data-bbox="834 1157 1484 1686"> <tbody> <tr> <td data-bbox="834 1157 1159 1241">Network Name field</td> <td data-bbox="1164 1157 1484 1241">Enter the name for the external network.</td> </tr> <tr> <td data-bbox="834 1247 1159 1331">Network IP Start field</td> <td data-bbox="1164 1247 1484 1331">Enter the starting floating IPv4 address.</td> </tr> <tr> <td data-bbox="834 1337 1159 1421">Network IP End field</td> <td data-bbox="1164 1337 1484 1421">Enter the ending floating IPv4 address.</td> </tr> <tr> <td data-bbox="834 1428 1159 1512">Network Gateway field</td> <td data-bbox="1164 1428 1484 1512">Enter the IPv4 address for the Gateway.</td> </tr> <tr> <td data-bbox="834 1518 1159 1602">DNS Server field</td> <td data-bbox="1164 1518 1484 1602">Enter the DNS server IPv4 address.</td> </tr> <tr> <td data-bbox="834 1608 1159 1686">Subnet</td> <td data-bbox="1164 1608 1484 1686">Enter the Subnet for External Network.</td> </tr> </tbody> </table>	Network Name field	Enter the name for the external network.	IP Start field	Enter the starting floating IPv4 address.	IP End field	Enter the ending floating IPv4 address.	Gateway field	Enter the IPv4 address for the Gateway.	DNS Server field	Enter the DNS server IPv4 address.	Segmentation ID field	Enter the segmentation ID.	Subnet	Enter the Subnet for Provider Network.	Network Name field	Enter the name for the external network.	Network IP Start field	Enter the starting floating IPv4 address.	Network IP End field	Enter the ending floating IPv4 address.	Network Gateway field	Enter the IPv4 address for the Gateway.	DNS Server field	Enter the DNS server IPv4 address.	Subnet	Enter the Subnet for External Network.
Network Name field	Enter the name for the external network.																										
IP Start field	Enter the starting floating IPv4 address.																										
IP End field	Enter the ending floating IPv4 address.																										
Gateway field	Enter the IPv4 address for the Gateway.																										
DNS Server field	Enter the DNS server IPv4 address.																										
Segmentation ID field	Enter the segmentation ID.																										
Subnet	Enter the Subnet for Provider Network.																										
Network Name field	Enter the name for the external network.																										
Network IP Start field	Enter the starting floating IPv4 address.																										
Network IP End field	Enter the ending floating IPv4 address.																										
Network Gateway field	Enter the IPv4 address for the Gateway.																										
DNS Server field	Enter the DNS server IPv4 address.																										
Subnet	Enter the Subnet for External Network.																										

Name	Description												
<p>TLS This optional section will only be visible once TLS is selected from Blueprint Initial Setup Page.</p>	<p>TLS has two options:</p> <ul style="list-style-type: none"> • External LB VIP FQDN - Text Field. • External LB VIP TLS - True/False. By default this option is false. 												
<p>SwiftStack optional section will be visible once SwiftStack is selected from Blueprint Initial Setup Page. SwiftStack is only supported with KeyStonev2 . If you select Keystonev3, swiftstack will not be available for configuration.</p>	<p>Following are the options that needs to be filled for SwiftStack:</p> <table border="1" data-bbox="873 541 1523 1155"> <tbody> <tr> <td data-bbox="873 541 1198 657">Cluster End Point</td> <td data-bbox="1198 541 1523 657">IP address of PAC (proxy-account-container) endpoint.</td> </tr> <tr> <td data-bbox="873 657 1198 741">Admin User</td> <td data-bbox="1198 657 1523 741">Admin user for swift to authenticate in keystone.</td> </tr> <tr> <td data-bbox="873 741 1198 888">Admin Tenant</td> <td data-bbox="1198 741 1523 888">The service tenant corresponding to the Account-Container used by Swiftstack.</td> </tr> <tr> <td data-bbox="873 888 1198 1077">Reseller Prefix</td> <td data-bbox="1198 888 1523 1077">Reseller_prefix as configured for Keysone Auth,AuthToken support in Swiftstack E.g KEY_</td> </tr> <tr> <td data-bbox="873 1077 1198 1155">Admin Password</td> <td data-bbox="1198 1077 1523 1155">swiftstack_admin_password</td> </tr> <tr> <td data-bbox="873 1155 1198 1165">Protocol</td> <td data-bbox="1198 1155 1523 1165">http or https. Protocol that swiftstack is running on top</td> </tr> </tbody> </table>	Cluster End Point	IP address of PAC (proxy-account-container) endpoint.	Admin User	Admin user for swift to authenticate in keystone.	Admin Tenant	The service tenant corresponding to the Account-Container used by Swiftstack.	Reseller Prefix	Reseller_prefix as configured for Keysone Auth,AuthToken support in Swiftstack E.g KEY_	Admin Password	swiftstack_admin_password	Protocol	http or https. Protocol that swiftstack is running on top
Cluster End Point	IP address of PAC (proxy-account-container) endpoint.												
Admin User	Admin user for swift to authenticate in keystone.												
Admin Tenant	The service tenant corresponding to the Account-Container used by Swiftstack.												
Reseller Prefix	Reseller_prefix as configured for Keysone Auth,AuthToken support in Swiftstack E.g KEY_												
Admin Password	swiftstack_admin_password												
Protocol	http or https. Protocol that swiftstack is running on top												

Note When tenant type ACI/VLAN is selected then ACI INFO tab is available in blueprint setup.

Note When ACI/VLAN is selected then Tor switch from initial setup is mandatory.

Name	Description
APIC Hosts field	Enter host input. Example: <ip1 host1>:[port] . max of 3, min of 1, not 2;
apic_username field	Enter a string format.
apic_password filed	Enter Password.
apic_system_id field	Enter input as string. Max length 8.
apic_resource_prefix field	Enter string max length 6.
apic_tep_address_pool field	Allowed only 10.0.0.0/16
multiclass_address_pool field	Allowed only 225.0.0.0/15
apic_pod_id field	Enter integer(1- 65535)

Name	Description
apic_installer_tenant field	Enter String, max length 32
apic_installer_vrf field	Enter String, max length 32
api_l3out_network field	Enter String, max length 32

Note When Tenant Type is VTS/VLAN then VTS tab is available in blueprint setup.

Name	Description
VTS Day0 (checkbox)	True or false default is false.
VTS User name	Enter as string does not contain special characters.
VTS Password	Enter password
VTS NCS IP	Enter IP Address format.
VTC SSH Username	Enter a string
VTC SHH Password	Enter password

Note If vts day0 is enabled then SSH username and SSH password is mandatory.

If SSH_username is input present then SSH password is mandatory vice-versa

Under the openstack setup tab, the Vim_admins tab will only be visible once Vim_admins is selected from the Optional Features & Services under the Blueprint InitialSetup tab.	<p>Following are the options that needs to be filled for Vim Admins:</p> <ul style="list-style-type: none"> • Username - Text Field • Password - Password field. Admin hash password should always start with \$6
--	---

Step 12

If **Syslog Export**, **NFVBENCH**, **ENABLE_ESC_PRIV** is selected in **Blueprint Initial Setup** Page then, **Services Setup** page will be enabled for User to view. Following are the options under Services Setup Tab:

Name	Description	
Syslog Export	Following are the options for Syslog Settings:	
	Remote Host	Enter Syslog IP Address.
	Protocol	Only UDP is supported.
	Facility	Defaults to local5.
	Severity	Defaults to debug.
	Clients	Defaults to ELK
	Port	Defaults to 514 but can be modified by the User.

NFVBENCH	<p>Enable checkbox which by default is false.</p> <p>Add ToR info connected to switch:</p> <ul style="list-style-type: none"> • Select a TOR Switch. Switch- (switch name) • Enter the port number. For Example: eth1/5 . VTEP VLANS (mandatory and needed only for VTS/VXLAN,); Enter 2 different VLANs for VLAN1 and VLAN2. • NIC Ports: INT1 & INT2 Optional input, enter the 2 port numbers of the 4-port 10G Intel NIC at the management node used for NFVBench.
ENABLE_ESC_PRIV	<p>Enable the checkbox to set it as True. By default it is False.</p>

Step 13 Click **Offline validation** button to initiate an offline validation of the Blueprint.

Step 14 Once the **Offline validation** is successful, **Save** option will be enabled for you which when clicked would redirect you to the **Blueprint Management Page**.

Downloading Blueprint

Before you begin

You must have atleast one blueprint (In any state Active/In-Active or In-progress), in the **Blueprint Management Page**.

Step 1 Log in to **CISCO VIM Insight**.

Step 2 In the navigation pane, expand the **Pre-Install Section**.

Step 3 Click **Blueprint Management**.

Step 4 Go-to **Download** for any Blueprint under Action title. (**Download Button** > **Downward Arrow** (with tooltip Preview & Download YAML).

Step 5 Click the **Download** icon.
A pop to view the Blueprint in the YAML format is displayed.

Step 6 Click the **Download** button at the bottom left of the pop-up window.
YAML is saved locally with the same name of the Blueprint.

Validating Blueprint

Step 1 Log in to **CISCO VIM Insight**.

Step 2 In the **Navigation** pane, expand the **Pre-Install Section**.

- Step 3** Click **Blueprint Creation**.
- Step 4** Upload an existing YAML, or create a **New Blueprint**.
Fill all the mandatory fields so that all Red Cross changes to **Green Tick**.
- Step 5** Enter the name of the Blueprint.
- Step 6** Click **Offline Validation**.
Only, if the Validation is successful, the Insight allows you to save the blueprint.

What to do next

If you see any errors, a hyperlink is created for those errors. Click the link to be navigated to the page where error has been encountered.

Managing Post Install Features

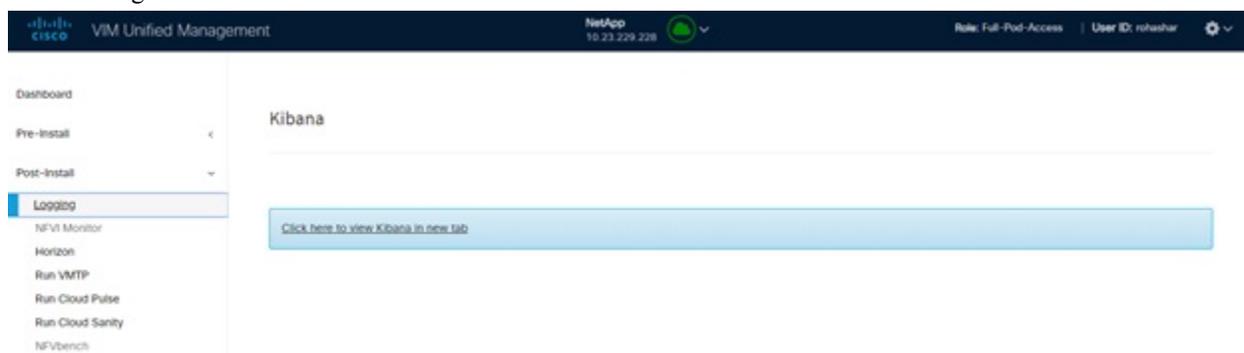
Cisco VIM provides an orchestration that helps in lifecycle management of a cloud. VIM is responsible for pod management activities which includes fixing both hardware and software issues with one-touch automation. VIM Insight provides the visualization of the stated goal. As a result, it integrates with POST install features that Cisco VIM offers through its Rest API. These features are enabled only if there is an active Blueprint deployment on the pod.

Monitoring the Pod

Cisco VIM uses EFK (Elasticsearch, Fluentd, and Kibana) to monitor the OpenStack services, by cross-launching the Kibana dashboard.

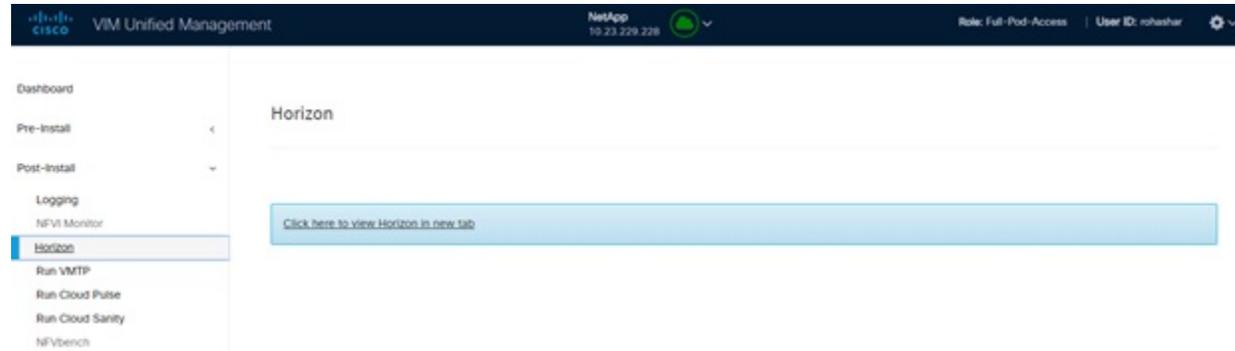
To cross launch Kibana, complete the following instructions:

- Step 1** In the navigation pane, click **Post-Install > Logging**.
- Step 2** Click **Click here to view Kibana in new tab**.
- Step 3** Enter the **Username** as Admin.
- Step 4** Enter the Kibana_PASSWORD password that is obtained from /root/installer-`<tagid>/openstack-configs/secrets.yaml` in the management node.



Cross Launching Horizon

Horizon is the canonical implementation of OpenStack's Dashboard, which provides a web-based user interface to OpenStack services including Nova, Swift and, Keystone.



Step 1 In the navigation pane, click **Post-Install > Horizon**.

Step 2 Click the link **Click here to view Horizon logs in new tab**. You will be redirected to Horizon landing page in a new tab.

Step 3 Enter the ADMIN_USER_PASSWORD obtained from `/root/installer-
<tagid>/openstack-configs/secrets.yaml` in the management node.

NFVI Monitoring

NFVI monitoring is the Cross launch browser same as Horizon. NFVI monitoring link is available in the post install only if the setupdata has NFVI Monitoring configuration during the cloud deployment. NFVI Monitoring checks the status of **Collector VM1 Info** and **Collector VM2 Info**.

Step 1 In the navigation pane, click **Post-Install > NFVI Monitoring**.

Step 2 Click the link **Click here to view NFVI monitoring**.

You will be redirected to NFVI Monitoring page.

Run VMTP

Run VMTP is divided in two sections:

- **Results for Auto Run:** This shows the results of VMTP which was run during the cloud deployment (Blueprint Installation).
- **Results for Manual Run:** Run the VMTP on demand. To run VMTP on demand, click **Run VMTP**.



Note If VMTP stage was skipped or has not-run during Blueprint Installation, this section of POST Install would be disabled for the user.

Run CloudPulse

In VIM 2.0 and later, we provide an integrated tool, called Cloud Pulse, that periodically checks the cloud services endpoint. The results of these tests are reflected under the Cloud Pulse link. Also, you can run these API endpoint tests on demand, and fetch the result of these tests by refreshing the table.

OpenStack CloudPulse tool is used to verify Cisco NFVI health. CloudPulse servers are installed in containers on all Cisco NFVI control nodes and CloudPulse clients are installed on the management node.

CloudPulse has two test sets: endpoint scenario (runs as a cron or manually) and operator test (run manually).

Following are the tests which are supported in CloudPulse:

Endpoint tests include

- cinder_endpoint
- glance_endpoint
- keystone_endpoint
- nova_endpoint
- neutron_endpoint

Operator tests include

- ceph_check
- docker_check
- galera_check
- node_check
- rabbitmq_check

CloudPulse

Cloudpulse Monitoring for: **Fixadent-BP**

Name	Result	State	Test Type	Created Date	Updated Date
neutron_endpoint	success	success	periodic	05/04/2018, 11:51:28	05/04/2018, 11:51:29
docker_check	All docker containers are ...	success	periodic	05/04/2018, 11:55:17	05/04/2018, 11:55:20
nova_endpoint	success	success	periodic	05/04/2018, 11:51:29	05/04/2018, 11:51:30
cinder_endpoint	success	success	periodic	05/04/2018, 11:55:20	05/04/2018, 11:55:27
keystone_endpoint	success	success	periodic	05/04/2018, 11:55:20	05/04/2018, 11:55:28
rabbitmq_check	Running Nodes : [rabbit...	success	periodic	05/04/2018, 11:55:20	05/04/2018, 11:55:27
galera_check	Active Nodes : 10.10.35...	success	periodic	05/04/2018, 11:55:22	05/04/2018, 11:55:25
glance_endpoint	success	success	periodic	05/04/2018, 11:55:28	05/04/2018, 11:55:28
neutron_endpoint	success	success	periodic	05/04/2018, 11:55:28	05/04/2018, 11:55:29
nova_endpoint	success	success	periodic	05/04/2018, 11:55:29	05/04/2018, 11:55:30

10 items per page

To run a cloud pulse test, choose a particular test from the dropdown and click **Run Test**. Once the test is in progress, Click (**Spin/refresh**) icon to fetch the latest result. This grid does not fetch the latest result automatically.

Run Cloud Sanity

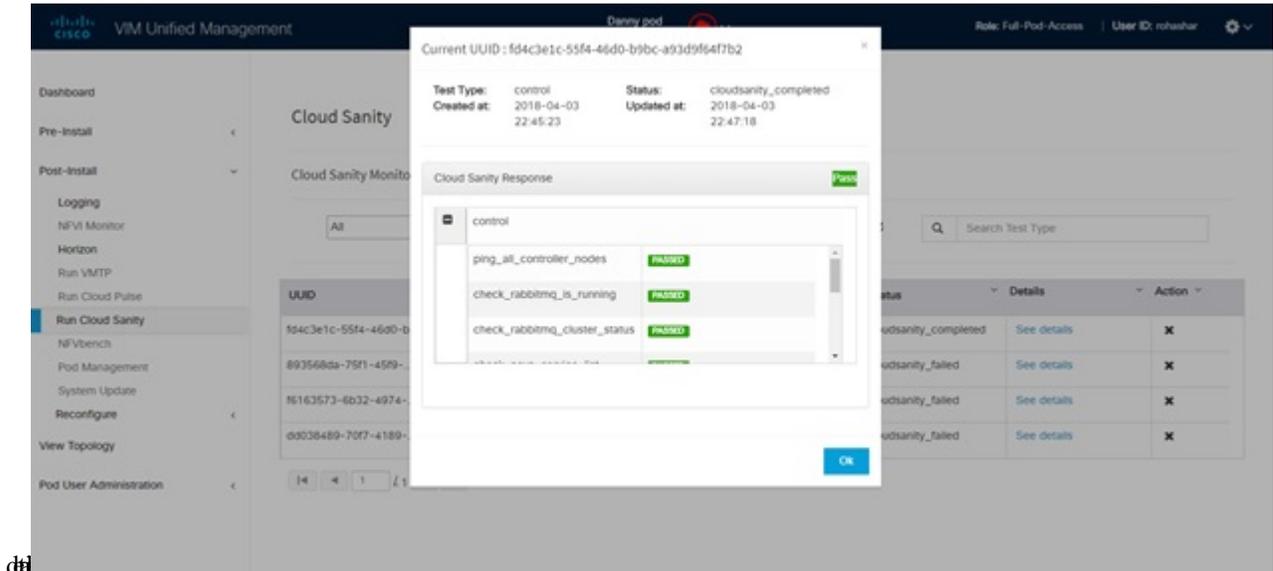
You can use the cloud sanity tool to test the Cisco NFVI pod and cloud infrastructure (host connectivity, basic mraiadb, rabbit, ceph cluster check, and RAID disks).

Following are the test available to run from insight.

- Control
- Compute
- Cephmon
- Cephosd
- Management
- All

Step 1 To run a Cloud sanity test choose a particular test from the dropdown.

Step 2 Click **Run Test** to proceed with the operation. Once the test is completed, click **See Details** for more



Run NFV Bench

You can **Run NFV Bench** for **B** and **C** series Pod, through Cisco VIM Insight. On a pod running with CVIM, choose a *NFVbench* link on the NAV-Menu.

You can run either fixed rate test or NDR/PDR test. As the settings and results for the test types differ, the options to run these tests are presented in two tabs, with its own settings and results.

NDR/PDR Test

- Step 1** Log in to **CISCO VIM Insight**.
- Step 2** In the Navigation pane, click **Post-Install** > Run NFV Bench.
- Step 3** Click NDR/PDR test and complete the following fields

Name	Description
Iteration Duration	Choose duration from 10 to 60 sec. Default is 20 sec.
Frame Size	Choose the correct frame size to run.
Run NDR/PDR test	Click Run NDR/PDR test . After, completion it displays each type of test with its own settings and results.

Reconfiguring CIMC Password Through Unified Management

Update the `cimc_password` in the CIMC-COMMON section, or the individual `cimc_password` for each server and then run the update password option.

To update a password, you have to follow the password rules:

- Must contain at least one lower-case letter.
- Must contain at least one upper-case letter.
- Must contain at least one digit between 0 to 9.
- One of these special characters !\$#@%^_+=*&
- Your password has to be 8 to 14 characters long.

Before you begin

You must have a C-series pod up and running with Cisco VIM to reconfigure CIMC password.



Note Reconfigure CIMC password section is disabled if the pod is in failed state as indicated by `ciscovim install-status`.

Step 1 Log in to **CISCO VIM Insight**.

Step 2 In the navigation pane, choose **Post-Install**

Step 3 Click **Reconfigure CIMC Password**.

Step 4 On the Reconfigure CIMC Password page of the Cisco VIM UM, complete the following fields:

Name	Description
CIMC_COMMON old Password	CIMC_COMMON old password field cannot be edited.
CIMC-COMMON new Password	Enter the CIMC-COMMON password. Password has to be alphanumeric according to the password rule.
Click Update	Old CIMC-COMMON password can be updated with new CIMC-COMMON password.



CHAPTER 9

Managing Pod Through Cisco VIM Unified Management

The following are the naming conventions used in the Cisco VIM UM

1. Super Administrator (UM Admin): User having access to UM Admin profile
2. POD Administrator: User having access to register a Pod in the system(Only UM can add new Pod Admin in the system)
3. Pod users (Normal users):
 - o All the users which are associated with the Pod. Full-pod-access: Role assigned to user which gives full access of a specific Pod(This has nothing to do with Pod Admins)

The following are the Key Points

- User who are UM admin or Pod admin but not associated with any Pod are not counted in UM admin dashboard user count section
- Only Pod Admins can register a new Pod
- Every Pod must a user with “Full-pod-Access” role.
- User cannot be revoked/delete if the users is the last user on the pod with “Full-Pod-Access” role.
- User cannot be delete if user is a Pod admin or UM admin.

The following topics tell you how to install and replace Cisco Virtual Infrastructure Manager (VIM) nodes using Cisco VIM Insight.

- [Managing Hardware, on page 183](#)
- [POD Management, on page 184](#)
- [Power Management, on page 191](#)
- [Managing Software, on page 194](#)
- [Pod User Administration, on page 208](#)

Managing Hardware

Management of your Cisco VIM pods includes adding, removing, or replacing the nodes.

In a pod, multiple nodes cannot be changed at the same time. For example, if you want to replace two control nodes, you must successfully complete the replacement of the first node before you begin to replace the second

node. Same restriction applies for addition and removal of storage nodes. Only, in case of Compute Nodes you can add or remove multiple nodes together. However, there must always be one active compute node in the pod at any given point. VNF manager stays active and monitors the compute nodes so that moving the VNFs accordingly as compute node management happens.



Note When you change a control, storage, or compute node in a Cisco VIM pod using Insight, it automatically updates the server and role in the active blueprint, as a result, your OpenStack deployment changes. When a node is removed from Cisco VIM, sensitive data may remain on the drives of the server. Administrator advice you to use Linux tools to wipe the storage server before using the same server for another purpose. The drives that are used by other application server must be wiped out before adding to Cisco VIM.

Searching Compute and Storage nodes

This functionality allows you to search the Compute and Storage nodes by server names only. The search result is generated or shows an empty grid if there are no results.

Figure 21: Search Storage Nodes

Server Na...	Rack Id	CIMC IP	CIMC Use...	CIMC Pas...	VIC Slot	Tor Info	Managem...	Managem...	Status	Actions	Logs
c38-storag...	RackA	172.26.229...		*****		(* c38-n9k...	192.168.38...		Active	View Logs	

Showing 1 of 1 items

POD Management

Cisco VIM allows the admin to perform pod life-cycle management from a hardware and software perspective. VIM provides the ability to power on/off compute node, add, remove or replace nodes based on the respective roles when the nodes of a given pod corrupts at times.

Figure 22: POD Management

The screenshot shows the Cisco VIM Unified Management interface. The top navigation bar includes the Cisco logo, 'VIM Unified Management', the user's name 'Danyal Pod' with a profile icon, and the role 'Role: Full-Pod-Access' and user ID 'User ID: Pravin Verma'. The left sidebar contains a menu with categories: Dashboard, Pre-Install (Blueprint Setup, Blueprint Management), Post-Install (Logging, NFVI Monitor, Horizon, Run VMTP, Run Cloud Pulse, NFVbench), Pod Management (System Update, Reconfigure, CIMC Password, Openstack Password, Openstack Config, Optional Services, View Topology), and Pod User Administration. The main content area is titled 'POD Management' and features a 'Summary' section with two cards: 'Node Summary' showing 9 Total nodes (3 Control, 3 Compute, 3 Storage) and 'IP Pool Summary' showing 11 Total nodes (2 Available). Below the summary is a table of nodes with columns: Server, Rack ID, CIMC IP, CIMC U..., CIMC P..., VIC Slot, Tor Info, DP Tor L..., SRIOV..., Manage..., Status, and Actions. The table lists three nodes, all with 'Active' status. The bottom of the table shows '1 of 3 of 3 items' and a pagination control for 5 items per page.

Pod Management page has two sections–

1. **Node Summary:** Node Summary section shows how many nodes are available and the detailed count of Control, Compute and Storage type.
2. **IP Pool Summary:** IP Pool Summary section shows the Total Pool Summary and the current available pool count.

The operations performed on the running pod are:

Replace Control Nodes: We do not support double fault scenarios, replacement of one controller at a time is supported.

Add Computes/Storage Nodes: N-computes nodes can be replaced simultaneously; however at any given point, at least one compute node has to be active.

Power On/ Off compute Nodes: You can Power On or Power Off compute node. At least one compute node must be powered on.

Remove Compute/Storage Nodes: You can add one node at a time, given that we run Ceph as a distributed storage offering.

Add Pool: You can increase pool size at any time.

Managing Storage Nodes

Before you add or remove a storage node, review the following guidelines for Managing Storage Nodes.

- **Required Number of Storage Nodes:** A Cisco VIM pod must have a minimum of three and a maximum of 20 storage nodes. If your pod has only two storage nodes, you cannot delete a storage node until you

add another storage node. If you have fewer than three storage nodes, you can add one node at a time until you get to 20 storage nodes.

- **Validation of Nodes:** When you add a storage node to a pod, Cisco VIM Insight validates that all the nodes in the pod meet the minimum requirements and are in active state. If you have a control or compute node in a faulty state, you must either correct, delete or replace that node before you can add a storage node.
- **Update Blueprint:** When you add or delete a storage node, Insight updates the blueprint for the Cisco VIM pod.
- **Storage Node Logs:** You can access the logs for each storage node from the link in the Log column on the **Storage Nodes** tab.

Adding Storage Node

Complete the following instructions to add a storage node:



Note You cannot add more than one storage node at a time.

Before you begin

- Remove the non-functional storage node from the pod. You can have maximum 20 storage nodes in a Cisco VIM pod.
- Ensure that the server for the new storage node is in powered state in OpenStack for C Series.

-
- Step 1** In the navigation pane, choose **Post-Install > Pod Management > Storage**.
- Step 2** Click on Add Storage node button on the Storage tab. A popup will open where you can provide information about the new Storage node.
- Step 3** For C Series, add the following details:
- **Server Name:** Name for the Storage Server to be added.
 - **Rack ID:** Enter the Rack ID. (Accepts String format).
 - **CIMC IP:** Enter the CIMC IP.
 - **CIMC User Name:** User name for the CIMC.
 - **CIMC Password:** Enter the password for the CIMC
 - **VIC Slot:** Enter the VIC Slot (Optional).
 - **ToR switch info:** Mandatory if ToR is configured as True
 - **Management IPv6:** Enter IPv6 Address.
- Step 4** For B Series, add the following details:
- **Server Name:** Name for the Storage Server to be added.

- **Rack ID:** Enter the Rack ID. (Accepts String format).
 - **Rack Unit ID:** Enter the Rack Unit ID.
 - **Management IPv6:** Enter IPv6 Address.
- Note** Cancel will discard the changes and popup will be closed

If all mandatory fields are filled in correctly then **Add Storage** button will be enabled.

Step 5 Click **Initiate Add Storage**. Add node initialized message will be displayed.

Step 6 To view logs, click **View logs** under Logs column.
The status of the POD will change to Active.

Step 7 Two kinds of failure may occur:

- **Add Node Pre-Failed:** When addition of node failed before the bare-metal stage (step 4) the Active Blueprint will be modified but the Node is not yet added in the Cloud. If you press **X** Icon, then Insight will delete the node information from the Blueprint and the state would be restored.
- **Add Node Post-Failed:** When addition of node failed after the bare-metal stage (step 4) the Active Blueprint will be modified and the node is registered in the cloud. If you press **X** Icon, then Insight will first delete the node from the Blueprint and then node removal from cloud would be initiated.

You can view the logs for this operation under **Logs** column.

Deleting Storage Node

You cannot delete more than one storage node at a time.

Step 1 In the Navigation pane, choose **Post-Install > POD Management > Storage**.

Step 2 Click **X** adjacent to the storage node you want to delete.

Step 3 **Node Removal Initiated successfully** message will be displayed.

To view logs, click **View logs** under logs column.

- If the Storage Node is deleted successfully, the storage node will be removed from the list under **Add/Remove storage Node**.
- In deletion failed, a new button **Clear Failed Nodes** will be displayed. Click **Clear Failed Nodes** to remove the node from cloud and Blueprint.

Managing Compute Nodes

Before you add or remove a compute node, review the following guidelines:

- **Required Number of Compute Nodes:** Cisco VIM pod must have a minimum of one compute node and a maximum of 61 compute nodes (with 3 ceph nodes). If your pod has only one compute node, you cannot delete that node until you add another compute node.

- **Update Blueprint:** When you add or remove a compute node, Insight updates the blueprint for the Cisco VIM pod.
- **Compute Node Logs:** You can access the Logs for each compute node from the link in the Log column on the Compute Nodes table.

Adding Compute Node

Add IP Pool

If all the existing pool size is already used, then you need to increase the pool size. On the Add compute or Add storage popup, Click **Expand Management IP pool** to add a new Pool.

The screenshot shows a dialog box titled "Expand Management IP pool". It contains the following fields and values:

- Subnet : 10.1.1.0/24
- Gateway : 10.1.1.9
- VLAN ID : 3333
- Management Node IP: IPv4 (selected), IPv6 (unselected)
- Existing IPv4 Pool: * 10.1.1.11 to 10.1.1.20, 10.1.1.21
- Add IPv4 Pool: * Enter New Management/Provision Pool

Complete the instructions, to add a compute node:

Before you begin

Ensure that the server for the new compute node is in powered state in OpenStack. You can add more than one compute node at a time.

Step 1 In the navigation pane, click **Post-Install > Pod Management > Compute**.

Step 2 Click **Add Compute Node** on the Compute tab a popup opens . Add the required information in the popup. To add another node click **Add Another Node** if you planned to add another compute node OR hit Initiate Add Compute if you so not plan to add any more compute node. If you hit “Add Another Node” button, the existing form will be emptied. You will have to fill the information for the new compute node and then repeat step 1. You may use Previous and Next button to navigate among different added node information.

Step 3 For C Series, add the following details:

- **Server Name:** Name for the Compute Server.
- **Rack ID:** Enter the Rack ID. (Accepts String format).
- **CIMC IP:** Enter the CIMC IP.
- **CIMC User Name:** User name for the CIMC.
- **CIMC Password:** Enter the password for the CIMC.

- **VIC Slot:** Enter the VIC Slot (Optional).
- **ToR switch info:** Mandatory if configured ToR is true.
- **DP ToR switch info:** Enter input as string format.
- **SRIVO ToR info :** Enter input as string format.
- **Management IPv6 :** Enter IPv6 Address.

Step 4 For B Series, add the following details:

- **Server Name:** Name for the Storage Server to be added.
- **Rack ID:** Enter the Rack ID. (Accepts String format).
- **Rack Unit ID:** Enter the Rack Unit ID.
- **Chassis ID:** Enter the Chassis ID. Range for Chassis ID is 1-24.
- **Blade ID:** Enter the Blade ID. Range for Blade ID is 1-8.
- **CIMC Password:** Enter the CIMC Password.
- **Management IPv6:** Enter IPv6 address.

If all mandatory fields are filled in correctly then click **Save**

Note Add Compute process can initiate multiple add of compute nodes. Fill in the mandatory fields to save new compute node or press cancel to exit message will be displayed.

Fields of Pod management will remain mandatory for user input based on setup-data.

Step 5 You may perform one among these steps mentioned below:

- Clicking **Cancel** displays the compute node information listed in the table and **Add Compute Node** button is enabled.
- If you feel you have filled in a wrong entry for the compute node information, click **Delete**. This will delete the entry from the table as this information is not added in the Blueprint.
- Click **Initiate Add Compute**, displays Add node initialized message.

Step 6 To view logs, click **View logs** under Logs column. The status of the POD will change to Active.

Step 7 Two kinds of failure may occur:

- **Add Node Pre-Failed:** When addition of node failed before the bare-metal stage (step 4) the Active Blueprint will be modified but the Node is not yet added in the Cloud. If you press **X** Icon, then Insight will delete the node information from the Blueprint and the state would be restored.
- **Add Node Post-Failed:** When addition of node failed after the bare-metal stage (step 4) the Active Blueprint will be modified and the node is registered in the cloud. If you press **X** Icon, then Insight will first delete the node from the Blueprint and then node removal from cloud would be initiated.

You can view the logs for this operation under **Logs** column.

Deleting Compute Node

Compute node is deleted due to a hardware failure. You can delete one compute node at a time.



Note If your pod has only one compute node, you cannot delete that node until you add another compute node.

-
- Step 1** In the navigation pane, choose **Post-Install > POD Management > Compute**.
- Step 2** Click **X** for the compute node to be deleted. To remove multiple compute nodes, choose the target compute nodes which is on the extreme left column, then click **Trash** Icon to remove multiple computes. Node Removal Initiated successfully message is displayed.
- Step 3** To view the Logs, click **View logs** under Logs column.
- If compute nodes are deleted successfully, you cannot view the compute node in the list under **Add or Remove Compute Node**.
 - If Compute Note is deleted, a new button **Clear Failed Nodes** is displayed.
- Step 4** Click **Clear Failed Nodes** to remove the node form Cloud and Blueprint.
-

Managing Control Nodes

Before you replace a control node, review the following guidelines:

- **Required Number of Control Nodes:** A Cisco VIM pod must have three control nodes and you can only replace one node at a time.
- **Validation of Nodes:** When you replace a control node, Cisco VIM Insight validates if all the other nodes in the pod meet the minimum requirements and are in active state. If you have a storage or a compute node in a faulty state, you must correct the faulty state or delete or replace that node before you can replace the control node.
- **Update Blueprint:** When you replace a control node, Insight updates the Active blueprint for the Cisco VIM pod.
- **Control Node Logs:** You can access the logs for each control node from the link in the **Logs** column of Control Nodes table.

Replacing Control Node

You can replace only one control node at a time.

- Step 1** In the navigation pane, click **Post-Install > Pod Management > Control**.
- Step 2** Click (Spin) icon. A confirmation pop-up appears, Click proceed to continue.
- Step 3** If you want to edit a specific control node before replace, click **Edit** to update the changes.
- Step 4** On success, **Replace Node Initiated** successfully message is displayed.

Step 5 You can view the logs in the **Logs** column on the Control Nodes table.

What to do next

If the replacement of the control node fails, do the following:

- Click the link in the Logs column.
- Check the logs to determine the cause of the failure.
- Correct the issue and attempt to replace the control node again.

Power Management

Compute node can be powered on or powered off from the Compute Tab in Pod Management section. There is a power button associated with each compute with information provided as tooltip when you hover on that icon.

Following are the steps to power on/off multiple compute node:

1. Click **Power** button located to the left of delete button.
2. Choose the compute nodes by selecting the check box, the corresponding power button gets enabled.

Power On a Compute Node

Following are the steps to power on the compute node:

1. Click the **Compute** tab.
2. In the Pod Management area, check the check box corresponding to the Compute node that you want to power on.



Note The **Power** button of a Compute node is enabled only after you select the Compute node.

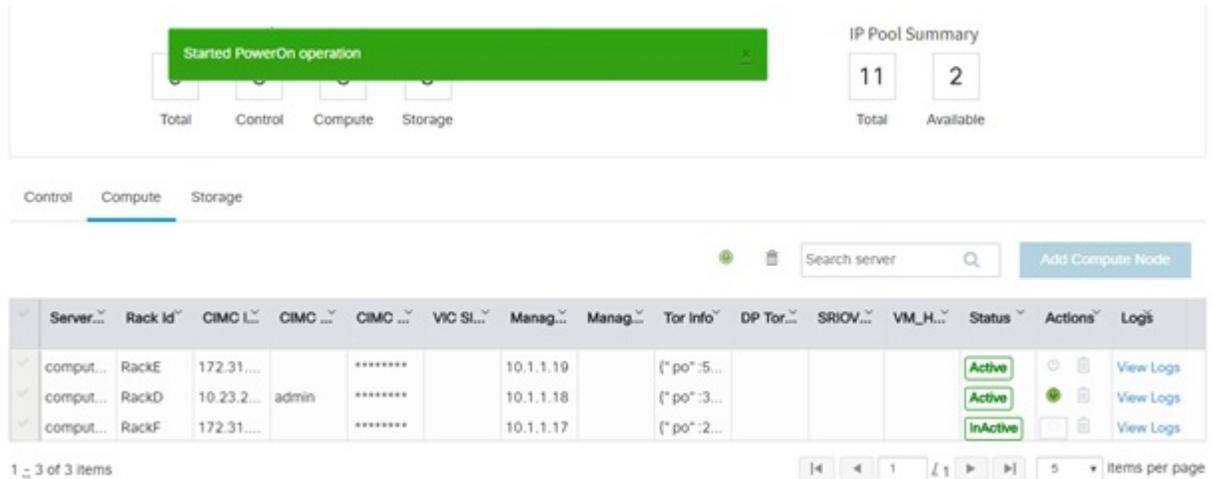
Figure 23: Powering On a Compute Node

Server	Rack Id	CIMC L	CIMC ...	CIMC ...	VIC SL	Manag...	Manag...	Tor Info	DP Tor...	SRIOV...	VM_H...	Status	Actions	Logs
comput...	RackE	172.31...		*****		10.1.1.19		{^ po":5...				Active		View Logs
comput...	RackD	10.23.2...	admin	*****		10.1.1.18		{^ po":3...				Active		View Logs
comput...	RackF	172.31...		*****		10.1.1.17		{^ po":2...				InActive		Power On, Click to Power Off

1 - 3 of 3 items

- Under the Actions column, click the **Power** button of the Compute node. It may take a few minutes for the Compute node to power on. The tooltip of the power button displays the status of the Compute node. Once the compute node is powered on, the Power button stops blinking and its color changes to green.

Figure 24: Power On Operation



The screenshot displays the Cisco VIM Unified Management interface. At the top, a green tooltip indicates "Started PowerOn operation". Below this, the IP Pool Summary shows 11 Total and 2 Available. The Compute tab is selected, and a table lists three compute nodes with their status (Active, Active, InActive).

Server...	Rack Id	CIMC L...	CIMC ...	CIMC ...	VIC SI...	Manag...	Manag...	Tor Info	DP Tor...	SRIOV...	VM_H...	Status	Actions	Logs
comput...	RackE	172.31....		*****		10.1.1.19		(*po"-5...				Active	View Logs	
comput...	RackD	10.23.2...	admin	*****		10.1.1.18		(*po"-3...				Active	View Logs	
comput...	RackF	172.31....		*****		10.1.1.17		(*po"-2...				InActive	View Logs	

You can add a Compute node only once a power on task is complete.

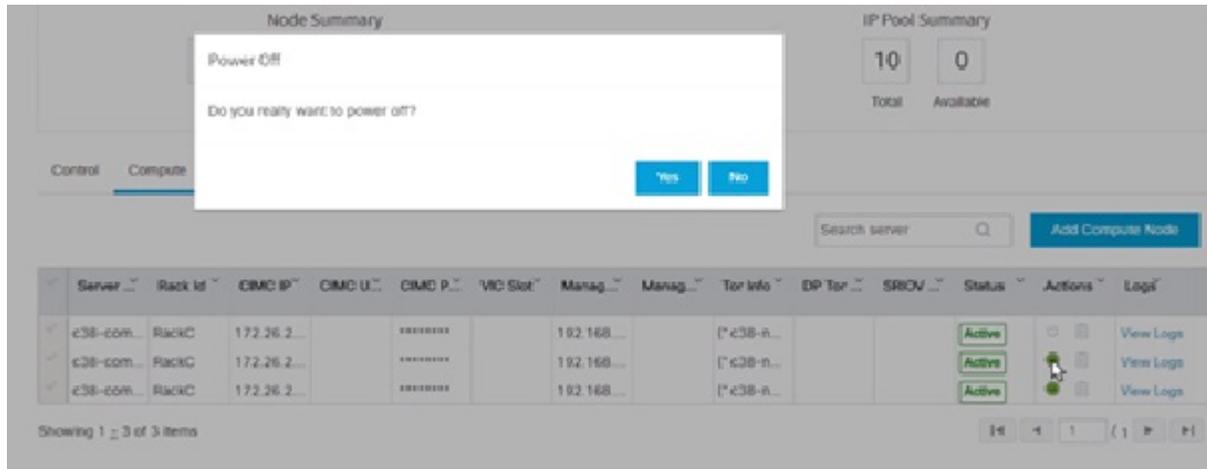
Powering Off Compute Node



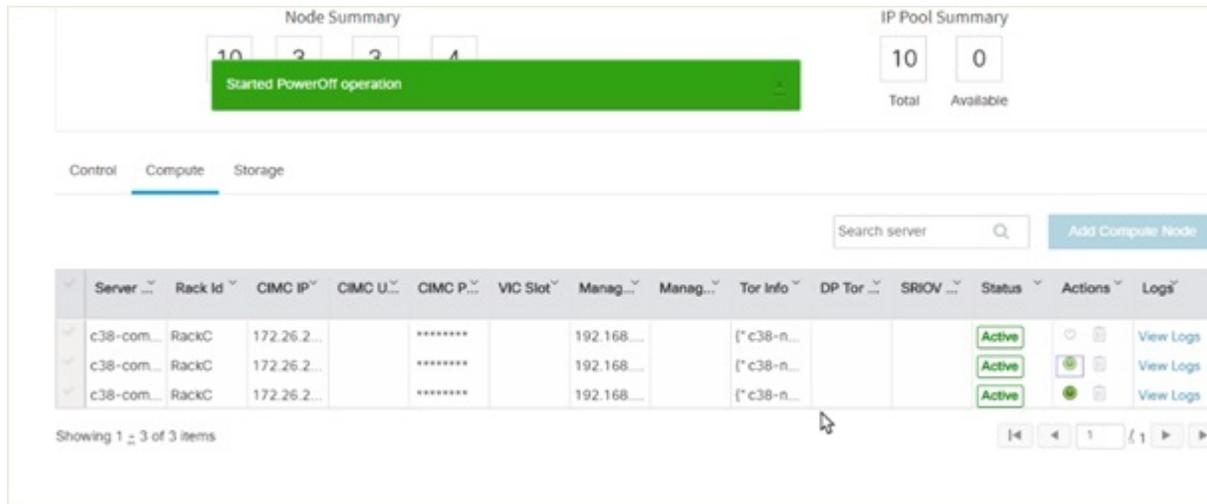
Note You cannot power off all the Compute nodes. There must be at least one Compute node that is in the On state.

Follow these steps to power off a Compute node:

- Click the **Compute** tab.
- In the Pod Management area, under the Actions column, click the **Power** button of the Compute node that you want to power off.



3. Click **Yes** in the confirmation dialog box.



It may take a few minutes for the Compute node to power off. The tooltip of the power button displays the status of the Compute node. Once the compute node is powered off, the Power button stops blinking and its color changes to grey.

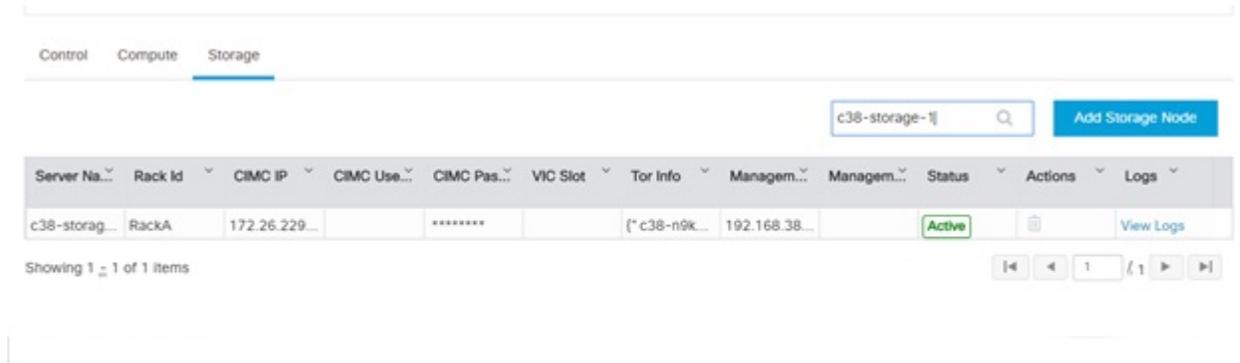


Note If there is only one compute node in the grid, and you try to power off it, a message *Last compute node can't be powered off* is displayed. Also, when you power off the last available compute node in the list of nodes, then the message *At least one compute node should be powered on* is displayed.

Searching Compute and Storage Nodes

This functionality allows you to search the Compute and Storage nodes by server names only. The search result is generated or shows an empty grid if there are no results.

Figure 25: Search Storage Nodes



Managing Software

Software management of your Cisco VIM pods includes software update, reconfigure of openstack services and password, etc.

VIM Software Update

As part of the lifecycle management of the cloud, VIM has the ability to bring in patches (bug fixes related to code, security, etc.), thereby providing cloud management facility from software point of view. Software update of the cloud is achieved by uploading a valid tar file, following initiation of a System Update form the Insight as follows:

-
- Step 1** In the Navigation pane, click **Post-Install > System Update**.
 - Step 2** Click **Browse** and select the valid tar file.
 - Step 3** Click **Open**.
 - Step 4** Click **Upload and Update**.
Update started Successfully message will be displayed.
 - Step 5** Update status will be shown as **ToUpdate**.
Click the hyperlink to view the reconfigure logs for install logs.
Reconfigure status will be available on the page or the dashboard under **POD Operation** details.
-

What to do next

System Update has been initiated message will be displayed. Logs front-ended by hyperlink will be in the section below in-front of **Update Logs** which shows the progress of the update. During the software update, all other pod management activities will be disabled. Post-update, normal cloud management will commence. Once update has completed you will see the status of update in the box below.

If log update fails, **Auto-RollBack** will be initiated automatically.

If log update is Successful, you will have two options to be performed:

1. **Commit**—To proceed with the update.
2. **RollBack**—To cancel the update.

If Auto-rollback fails during software update fails through Insight UI, it is advised that the administrator contact Cisco TAC for help. Do not re-try the update or delete the new or the old installer workspace.

Reconfigure Openstack Passwords

There are two options to regenerate the passwords:

- **Regenerate all passwords:** Click **Regenerate all passwords** checkbox and click **Set Password**. This will automatically regenerate all passwords in alphanumeric format.
- **Regenerate single or more password:** This will set a specific password by doing an inline edit for any service like Horizon's ADMIN_USER_PASSWORD. Double click on the field under Password and enter the password to enable **Set Password** button.

The screenshot shows the 'Openstack Password' configuration page in the Cisco VIM Unified Management interface. The page has a dark blue header with the Cisco logo and 'VIM Unified Management' text. A navigation sidebar on the left lists various management tasks. The main content area is titled 'Openstack Password' and features a 'Regenerate all Passwords' checkbox. Below this is a table with two columns: 'Password Name' and 'Password'. The table contains the following entries:

Password Name	Password
HEAT_KEYSTONE_PASSWORD	
CINDER_KEYSTONE_PASSWORD	
RABBITMQ_ERLANG_COOKIE	
METADATA_PROXY_SHARED_SECRET	
WSREP_PASSWORD	
ETCD_ROOT_PASSWORD	
HEAT_DB_PASSWORD	
CINDER_DB_PASSWORD	
KEYSTONE_DB_PASSWORD	
NOVA_DB_PASSWORD	

At the bottom of the table, there are 'Set Password' and 'Clear' buttons. To the right of the table, there is a 'Status Of Password Reconfigure' section with fields for 'Created At', 'Updated At', 'Status', 'Update Logs', and 'Actions'. The footer of the page contains copyright information: '©2018 Cisco and/or its affiliates. All rights reserved. Cisco VIM Unified Management Version: 2.2.2'.

During the reconfiguration of password, all other pod management activities will be disabled. Post-update, normal cloud management will commence. If the reconfigure of the password fails, all subsequent pod management operations will be blocked. It is advised to contact Cisco TAC to resolve the situation through CLI.

Reconfigure OpenStack Services, TLS Certificates, and ELK Configurations

Cisco VIM supports the reconfiguration of OpenStack log level services, TLS certificates, and ELK configuration. Following are the steps to reconfigure the OpenStack and other services:

Step 1 In the navigation pane, click **Post-Install > Reconfigure Openstack Config**.

- Step 2** Click the specific item that you want to change and update. For example: to update the TLS certificate click the path to the certificate location.
- Step 3** Enter **Set Config** to commence the process.

What to do next

During the reconfiguration process, all other pod management activities are disabled. Post-update, normal cloud management commences. If reconfigure of OpenStack Services fails, all subsequent pod management operations are blocked. Contact, Cisco TAC to resolve the situation through CLI.

Reconfiguring CIMC Password through Unified Management

Cisco VIM allows you to Update the `cimc_password` in the CIMC-COMMON section, and/or the individual `cimc_password` for each server and then run the update password option.

You need to match the following Password rule to update the Password:

- Must contain at least one lower case letter.
- Must contain at least one upper case letter.
- Must contain at least one digit between 0 to 9.
- One of these special characters `!$#@%^_+=*&`
- Your password has to be 8 to 14 characters long.

Before you begin

You must have a C-series pod up and running with Cisco VIM to reconfigure CIMC password.



Note Reconfigure CIMC password section will be disabled if the pod is in failed state as indicated by `ciscovim install-status`.

- Step 1** Log-in to **CISCO VIM Insight**.
- Step 2** In the navigation pane, select **Post-Install**.
- Step 3** Click **Reconfigure CIMC Password**.
- Step 4** You can reconfigure the CIMC Password at global level by adding new `CIMC_COMMON` Password or to reconfigure CIMC Password for individual servers double click the server password you want to edit.
- Step 5** Click **Reconfigure** to initiate reconfigure process.
-

Reconfigure Optional Services

Cisco VIM offers optional services such as heat, migration to Keystone v3, NFVBench, NFVIMON, etc, that can be enabled post-pod deployment. These services can be enabled in one-shot or selectively.

Listed below are the steps to enable optional services:

Step 1 In the Navigation pane, click **Post-Install > Reconfigure Optional Services**.

Step 2 Choose the right services and update the fields with the right values.

Step 3 Click **Offline validation**. Once offline validation is successful.

Step 4 Click **Reconfigure** to commence the process.

During the reconfiguration process, all other pod management activities will be disabled. Post-update, normal cloud management will commence.

If reconfigured OpenStack Services fail, all subsequent pod management operations are blocked. Contact Cisco TAC to resolve the situation through CLI.

Note All reconfigure operation feature contains repeated deployment true or false.

- Repeated re-deployment true - Feature can be re-deployed again.
- Repeated re-deployment false- Deployment of feature allowed only once.

Deployment Status :

Optional Features	Repeated re-deployment Option
APICINFO	True
EXTERNAL_LB_VIP_FQDN	False
EXTERNAL_LB_VIP_TLS	False
INSTALL_MODE	True
HTTP_PROXY & HTTPS_PROXY	True
LDAP	True
NETWORKING	True
NFVBENCH	False
NFVIMON	False
PODNAME	False
PROVIDER_VLAN_RANGES	True
SWIFTSTACK	True
SYSLOG_EXPORT_SETTINGS	False
TENANT_VLAN_RANGES	True
TORSWITCHINFO	False
VIM_ ADMINS	True

Optional Features	Repeated re-deployment Option
VMTP	False
VTS_PARAMETERS	False
AUTOBACKUP	True
Heat	False
Keystone v3	False

Reconfiguring Optional Features Through Unified Management

Step 1 Log-in to Cisco VIM UM.

Step 2 In the **Navigation** pane, expand the **Post-Install Section**.

Step 3 Click **Reconfiguring Optional Feature through UM**.

Step 4 On the **Reconfiguring Optional Feature through UM** page of the Cisco VIM UM, complete the following fields:

Name	Description
Heat check box	<ul style="list-style-type: none"> • Enable Heat. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process..
Keystone v3 check box	<ul style="list-style-type: none"> • Enable Keystone v3. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
ENABLE_ESC_PRIV	<ul style="list-style-type: none"> • Enable ENABLE_ESC_PRIV . • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
Autobackup check box	<ul style="list-style-type: none"> • Enable/Disable Autobackup. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.

Name	Description
External LB VIP TLS check box	<ul style="list-style-type: none"> • Enable External LB VIP TLS. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
External LB VIP FQDN check box	<ul style="list-style-type: none"> • Enter Input as a string. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
Pod Name	<ul style="list-style-type: none"> • Enter Input as a string. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
Tenant Vlan Ranges	<ul style="list-style-type: none"> • Augment tenant vlan ranges input. For Example: 3310:3315. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
Provider VLAN Ranges	<ul style="list-style-type: none"> • Enter input to tenant vlan ranges. For Example: 3310:3315. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.
Install Mode	<ul style="list-style-type: none"> • Select Connected or Disconnected, any one form the drop-down list. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process.

Name	Description	
Syslog Export Settings	Following are the options for Skylog Settings:	
	Remote Host	Enter Syslog IP Address.
	Facility	Defaults to local5
	Severity	Defaults to debug
	Clients	Defaults to ELK
	Port	Defaults to 514 but is modified by the User.
	Protocol	Supports only UDP
Configure ToR checkbox	True or False . Default is false.	

Name	Description		
<p>ToR Switch Information</p>	<p>Click + to add information for ToR Switch.</p>		
	<table border="1"> <thead> <tr> <th data-bbox="902 354 1214 388">Name</th> <th data-bbox="1221 354 1521 388">Description</th> </tr> </thead> </table>	Name	Description
	Name	Description	
	<p>Name</p>	<p>ToR switch name.</p>	
	<p>Username</p>	<p>ToR switch username.</p>	
	<p>Password</p>	<p>ToR switch Password.</p>	
	<p>SSH IP</p>	<p>ToR switch SSH IP Address.</p>	
	<p>SSN Num</p>	<p>ToR switch ssn num. output of show license host-id.</p>	
	<p>VPC Peer Keepalive</p>	<p>Peer Management IP. You need not define if there is no peer.</p>	
	<p>VPC Domain</p>	<p>Need not define if there is no peer.</p>	
	<p>VPC Peer port</p>	<p>Interface for vpc peer ports.</p>	
	<p>VPC Peer VLAN Info</p>	<p>vlan ids for vpc peer ports (optional).</p>	
	<p>BR Management Port Info</p>	<p>Management interface of the build node.</p>	
<p>BR Management PO Info</p>	<p>Port channel number for the management interface of the build node.</p>		
<p>Click Save</p> <ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 			

Note When setup data is ACI VLAN with TOR then reconfigure options are:

<p>TORSwitch Information mandatory table if you want to enter ToR information</p>	<p>Click + to add information for ToR Switch.</p> <table border="1" data-bbox="862 283 1482 592"> <thead> <tr> <th data-bbox="862 283 1175 338">Name</th> <th data-bbox="1175 283 1482 338">Description</th> </tr> </thead> <tbody> <tr> <td data-bbox="862 338 1175 394">Host Name</td> <td data-bbox="1175 338 1482 394">ToR switch name.</td> </tr> <tr> <td data-bbox="862 394 1175 451">VPC Peer Keepalive</td> <td data-bbox="1175 394 1482 451">Peer Management IP.</td> </tr> <tr> <td data-bbox="862 451 1175 508">VPC Domain</td> <td data-bbox="1175 451 1482 508">Do not define if there is no</td> </tr> <tr> <td data-bbox="862 508 1175 592">Node ID</td> <td data-bbox="1175 508 1482 592">Integer, unique across all switches</td> </tr> </tbody> </table> <p>Click Save</p> <ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 	Name	Description	Host Name	ToR switch name.	VPC Peer Keepalive	Peer Management IP.	VPC Domain	Do not define if there is no	Node ID	Integer, unique across all switches
Name	Description										
Host Name	ToR switch name.										
VPC Peer Keepalive	Peer Management IP.										
VPC Domain	Do not define if there is no										
Node ID	Integer, unique across all switches										
<p>NFV Bench</p>	<p>Enable check box which by default is false.</p> <p>Add Tor info connected to switch:</p> <ul style="list-style-type: none"> • Select a TOR Switch and Enter the Switch name. • Enter the port number. For example: eth1/5 • NIC Ports: INT1 and INT2 optional input, enter the 2 port numbers of the 4-port 10G Intel NIC at the management node used for NFVBench. • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. <p>Note If ToR is already present in Setup-data or already deployed. Then no need add Tor info, by default ToR info switchname is mapped in NFV bench.</p>										

<p>Swiftstack</p> <p>SwiftStack is only supported with Keystone v2. If you select Keystone v3, swiftstack will not be available for configuration.</p>	<p>Cluster End Point</p>	<p>IP address of PAC (proxy-account-container) endpoint.</p>
	<p>Admin User</p>	<p>Admin user for swift to authenticate in keystone.</p>
	<p>Admin Tenant</p>	<p>The service tenant corresponding to the Account-Container used by Swiftstack.</p>
	<p>Reseller Prefix</p>	<p>Reseller_prefix as configured for Keystone Auth,AuthToken support in Swiftstack E.g KEY_</p>
	<p>Admin Password</p>	<p>swiftstack_admin_password</p>
	<p>Protocol drop-down list</p>	<p>http or https</p>
	<ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 	

LDAP with Keystone v3	Domain Name field	Enter the Domain name.
	Object Class for Users field	Enter a string as input.
	Object Class for Groups	Enter a string.
	Domain Name Tree for Users	Enter a string.
	Domain Name Tree for Groups field	Enter a string.
	Suffix for Domain Name field	Enter a string.
	URL field	Enter a URL with port number.
	Domain Name for Bind User field	Enter a string.
	Password field	Enter Password as string format.
	User Filter	Enter filter name as string.
	User ID Attribute	Enter a string.
	User Name Attribute	Enter a string.
	User Mail Attribute	Enter a string.
Group Name Attribute	Enter a string.	
<ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 		

<p>NFV Monitoring</p>	<p>Followings are the field values for NFV Monitoring:</p> <table border="1"> <tr> <td data-bbox="901 283 1214 338">Master Admin IP field.</td> <td data-bbox="1219 283 1518 338">Enter Input as IP format.</td> </tr> <tr> <td data-bbox="901 344 1214 428">Collector Management IP field</td> <td data-bbox="1219 344 1518 428">Enter Input as IP format.</td> </tr> <tr> <td data-bbox="901 434 1214 478">Collector VM1 info</td> <td data-bbox="1219 434 1518 478"></td> </tr> <tr> <td data-bbox="901 485 1214 529">Host Name field</td> <td data-bbox="1219 485 1518 529">Enter Host Name as a string.</td> </tr> <tr> <td data-bbox="901 535 1214 590">CCUSER password field</td> <td data-bbox="1219 535 1518 590">Enter Password.</td> </tr> <tr> <td data-bbox="901 596 1214 640">Password field</td> <td data-bbox="1219 596 1518 640">Enter password.</td> </tr> <tr> <td data-bbox="901 646 1214 690">Admin IP field</td> <td data-bbox="1219 646 1518 690">Enter Input as IP format.</td> </tr> <tr> <td data-bbox="901 697 1214 741">Management IP field</td> <td data-bbox="1219 697 1518 741">Enter Input as IP format.</td> </tr> <tr> <td data-bbox="901 747 1214 791">Collector VM2 info</td> <td data-bbox="1219 747 1518 791"></td> </tr> <tr> <td data-bbox="901 798 1214 842">Host Namefield</td> <td data-bbox="1219 798 1518 842">Enter a string.</td> </tr> <tr> <td data-bbox="901 848 1214 892">CCUSER field</td> <td data-bbox="1219 848 1518 892">Enter Password.</td> </tr> <tr> <td data-bbox="901 898 1214 942">Management IP field</td> <td data-bbox="1219 898 1518 942">Enter Input as IP format.</td> </tr> <tr> <td data-bbox="901 949 1214 993">Dispatcher</td> <td data-bbox="1219 949 1518 993"></td> </tr> <tr> <td data-bbox="901 999 1214 1083">Rabbit MQ Username Field</td> <td data-bbox="1219 999 1518 1083">Enter a string.</td> </tr> </table> <ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 	Master Admin IP field.	Enter Input as IP format.	Collector Management IP field	Enter Input as IP format.	Collector VM1 info		Host Name field	Enter Host Name as a string.	CCUSER password field	Enter Password.	Password field	Enter password.	Admin IP field	Enter Input as IP format.	Management IP field	Enter Input as IP format.	Collector VM2 info		Host Namefield	Enter a string.	CCUSER field	Enter Password.	Management IP field	Enter Input as IP format.	Dispatcher		Rabbit MQ Username Field	Enter a string.
Master Admin IP field.	Enter Input as IP format.																												
Collector Management IP field	Enter Input as IP format.																												
Collector VM1 info																													
Host Name field	Enter Host Name as a string.																												
CCUSER password field	Enter Password.																												
Password field	Enter password.																												
Admin IP field	Enter Input as IP format.																												
Management IP field	Enter Input as IP format.																												
Collector VM2 info																													
Host Namefield	Enter a string.																												
CCUSER field	Enter Password.																												
Management IP field	Enter Input as IP format.																												
Dispatcher																													
Rabbit MQ Username Field	Enter a string.																												
<p>VTS Parameter</p>	<p>Following are the fields to reconfigure for VTS parameters</p> <table border="1"> <tr> <td data-bbox="901 1367 1214 1411">VTC SSH Username field.</td> <td data-bbox="1219 1367 1518 1411">Enter the string.</td> </tr> <tr> <td data-bbox="901 1417 1214 1461">VTC SSH Username field.</td> <td data-bbox="1219 1417 1518 1461">Enter the password.</td> </tr> </table> <ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 	VTC SSH Username field.	Enter the string.	VTC SSH Username field.	Enter the password.																								
VTC SSH Username field.	Enter the string.																												
VTC SSH Username field.	Enter the password.																												

VMTP	<p>Check one of the check boxes to specify a VMTP network:</p> <ul style="list-style-type: none"> • Provider Network • External Network <p>For the Provider Network complete the following:</p>	
	Network Name field.	Enter the name for the external network.
	IP Start field.	Enter the starting floating IPv4 address.
	IP End field.	Enter the ending floating IPv4 address.
	Gateway field	Enter the IPv4 address for the Gateway.
	DNS Server field.	Enter the DNS server IPv4 address.
	Segmentation ID field.	Enter the segmentation ID.
	Subnet	Enter the Subnet for Provider Network.
	<p>For External Network fill in the following details:</p>	
	Network Name field.	Enter the name for the external network.
	Network IP Start field.	Enter the starting floating IPv4 address.
	Network IP End field.	Enter the ending floating IPv4 address.
	Network Gateway field	Enter the IPv4 address for the Gateway.
	DNS Server field.	Enter the DNS server IPv4 address.
	Subnet	Enter the Subnet for External Network.
<ul style="list-style-type: none"> • Click Offline Validation . • When Offline Validation is successful, click Reconfigure to commence the process. 		

<p>Networking</p> <p>In Reconfigure optional services networking, you can reconfigure IP tables, or add http_proxy/https_proxy.</p>	<p>To reconfigure networking, update the relevant information:</p> <table border="1" data-bbox="901 283 1521 688"> <tr> <td data-bbox="901 283 1214 415">IP Tables</td> <td data-bbox="1219 283 1521 415">Click Add(+) to add a table. Enter input as subnet format. E.g. 12.1.0.1/2</td> </tr> <tr> <td data-bbox="901 415 1214 548">http_proxy_server</td> <td data-bbox="1219 415 1521 548">Enter HTTP_PROXY_SERVER E.g. <a.b.c.d:port></td> </tr> <tr> <td data-bbox="901 548 1214 688">https_proxy_server</td> <td data-bbox="1219 548 1521 688">Enter HTTP_PROXY_SERVER E.g. <a.b.c.d:port></td> </tr> </table> <ul style="list-style-type: none"> • Click Save. • Click Offline Validation. • When Offline Validation is successful, click Reconfigure to commence the process. 	IP Tables	Click Add(+) to add a table. Enter input as subnet format. E.g. 12.1.0.1/2	http_proxy_server	Enter HTTP_PROXY_SERVER E.g. <a.b.c.d:port>	https_proxy_server	Enter HTTP_PROXY_SERVER E.g. <a.b.c.d:port>
IP Tables	Click Add(+) to add a table. Enter input as subnet format. E.g. 12.1.0.1/2						
http_proxy_server	Enter HTTP_PROXY_SERVER E.g. <a.b.c.d:port>						
https_proxy_server	Enter HTTP_PROXY_SERVER E.g. <a.b.c.d:port>						
<p>APICINFO</p> <p>Note Reconfigure optional services only APIC hosts can be reconfigure.</p>	<p>To reconfigure APICINFO, follow the process:</p> <ul style="list-style-type: none"> • Enter input for APIC hosts format. <ip1 host1>:[port] or eg.12.1.0.12 • Click Save. • Click Offline Validation. • When Offline Validation is successful, click Reconfigure to commence the process. <p>Note APIC hosts can be reconfigure minimum 1 host and max 3 but not 2 hosts.</p>						
<p>Vim_admins</p>	<p>To reconfigure vim_admins, follow the process:</p> <ul style="list-style-type: none"> • To add a new root user, Click + and add the Username and admin hash password (Starting with \$6). • To remove the existing user, Click -. • When Offline Validation is successful, click Reconfigure to commence the process. 						

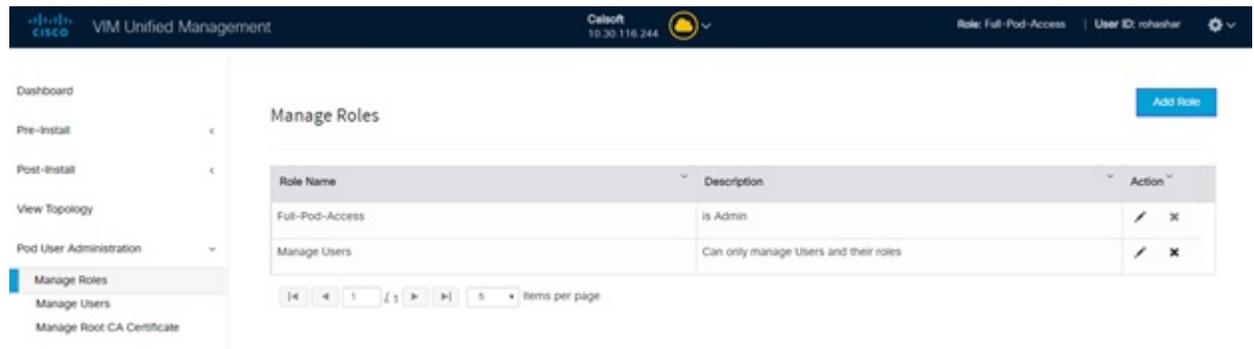
Pod User Administration

Cisco VIM UM offers Users (Pod Admins or Pod Users) to manage Users and roles that are associated with them.

Managing Roles

User can create multiple Roles and assign them to other pod users. System has a default role that is named as Full-Pod-Access which is assigned to the person who registers the Pod.

Manage Roles



Step 1 Click **Login as POD User**.

Step 2 Navigate to **Pod User Administration** and click **Manage Roles**. By default you see full-pod-access role in the table.

Step 3 Click **Add New Role** to create a new role.

Step 4 Complete the following fields in the **Add Roles** page in Cisco VIM UM:

Field Name	Field Description
Role	Enter the name of the role.
Description	Enter the description of the role.
Permission	Check the Permission check box to select the permission.
Click Save .	Once the Blueprint is in Active state all the permissions are same for C-series and B-series Pods other than Reconfigure CIMC Password which is missing for B-series Pod.

Note Permissions are divided in the granular level where viewing **Dashboard** is the default role that is implicitly added while creating a role.

Note Permissions are divided in the granular level where viewing **Dashboard** is the default role that is implicitly added while creating a role.

Managing Users

This section allows you to add the users. It shows all the users associated with the Pod. You can check the online status of all the user. Click **Refresh** on upper right corner to check the status.

To add a new user:

- Step 1** Click **Login as POD User**.
- Step 2** Navigate to **POD User Administration** and click **Manage Users**.
- Step 3** Click **Add Users** to add a new user.
- Step 4** Complete the following fields in the **Add Users** pane of the Cisco VIM Insight:

Field Name	Field Description
Email ID	Enter the Email ID of the User.
User Name	Enter the User Name if the User is new. If the User is already registered to the Insight the User-Name gets auto-populated.
Role	Select the Role from the drop-down list.

- Step 5** Click **Save** Once the Blueprint is in Active state all the permissions are same for C-series and B-series Pods other than Reconfigure CIMC Password which is missing for B-series Pod.

Revoke Users

User with Full-Pod-Access or Manage Users permission can revoke other users from the specific Pod.

To revoke users:

- Step 1** Click **Undo** icon. A confirmation pop up will appear.
- Step 2** Click **Proceed** to continue.

Note Self revoke is not permitted. After revoking the another user, if the user is not associated with any other pod then the revoked user will be auto deleted from the system.

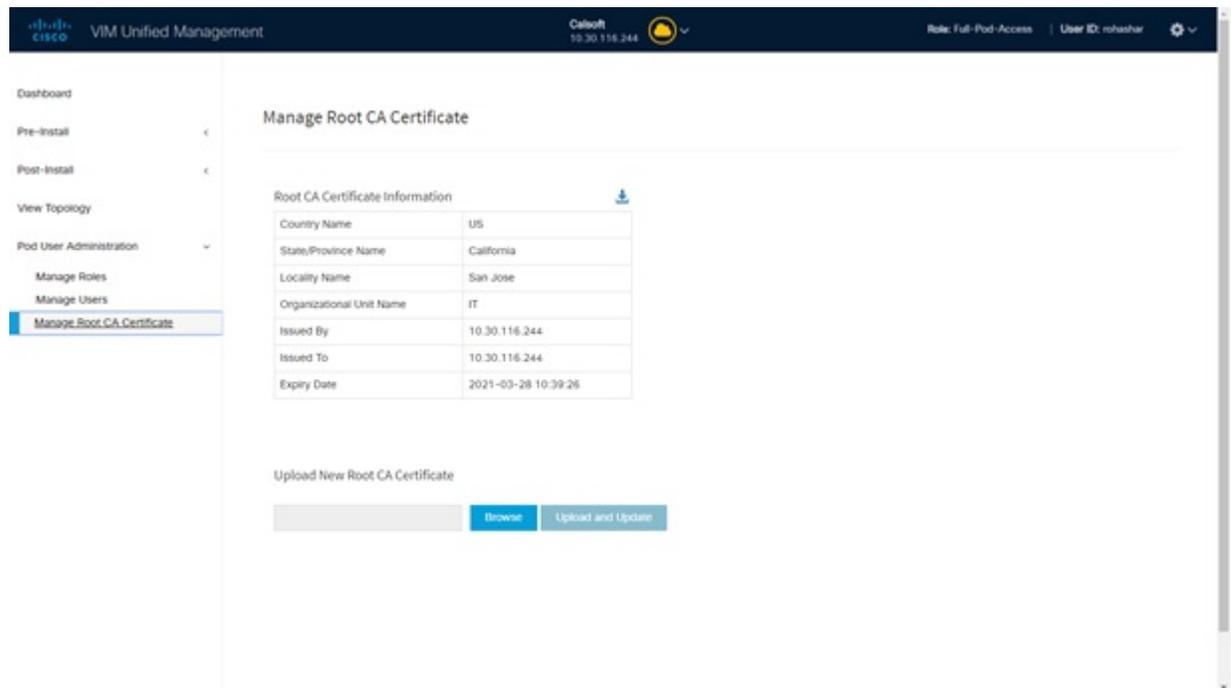
Edit Users

User with Full-Pod-Access or Manage Users permission can edit other user's permission for that specific Pod.
To edit user's permission

- Step 1** Click **Edit** icon.
- Step 2** Update the permission.
- Step 3** Click **Save**. The Grid will get refreshed automatically.

Managing Root CA Certificate

You can update the CA Certificate during the registration of the POD. Once, logged in as POD User and if you have the permission to update the certificate you can view under POD User Administration>> Manage Root CA Certificate.



The screenshot shows the 'Manage Root CA Certificate' page in the Cisco VIM Unified Management interface. The page includes a sidebar with navigation options and a main content area with the following details:

Root CA Certificate Information	
Country Name	US
State/Province Name	California
Locality Name	San Jose
Organizational Unit Name	IT
Issued By	10.30.116.244
Issued To	10.30.116.244
Expiry Date	2021-03-28 10:39:26

Below the table, there is an 'Upload New Root CA Certificate' section with a file input field, a 'Browse' button, and an 'Upload and Update' button.

To update the Certificate:

- Step 1** Click **Login as POD User**

Step 2 Navigate to **POD User Administration>>Manage Root CA certificate**.

Step 3 Click **Browse** and select the certificate that you want to upload.

Step 4 Click **Upload**.

- If the certificate is Invalid, and does not matches with the certificate on the management node located at (var/www/mercury/mercury-ca.crt) then Insight reverts the certificate which was working previously.
- If the Certificate is valid, Insight runs a management node health check and then update the certificate with the latest one.

Note The CA Certificate which is uploaded should be same as the one which is in the management node.



CHAPTER 10

Day 2 Operations of Cisco VIM Insight

The following topic guides you the details about the Day 2 Operations of Cisco VIM Insight.

- [Shutting Down Cisco VIM Unified Management, on page 213](#)
- [Restarting Cisco VIM Unified Management, on page 214](#)
- [Restoring VIM Insight, on page 214](#)
- [Reconfiguring VIM Unified Management , on page 216](#)
- [Reconfiguring Insight MySQL Database Password, on page 221](#)
- [Update VIM Insight , on page 226](#)
- [Rollback VIM Unified Management, on page 229](#)
- [Commit VIM Unified Management, on page 230](#)

Shutting Down Cisco VIM Unified Management

To stop the Cisco VIM Insight Container services, shut down Cisco UCS VIM Insight by running the `systemctl stop service` command.

Step 1 Log in to a server in which the Insight container is running.

Step 2 Stop the Insight service by running the following command from the Shell window:

```
systemctl stop docker-insight
```

a) Check the status of Insight Container by running the following command: `docker ps -a | grep insight`.

```
STATUS
Up 6 seconds
```

b) Check the status of the service by running the following command:

```
systemctl status docker-insight
```

The following information is displayed

```
Docker-insight.service - Insight Docker Service
Loaded: loaded (/usr/lib/systemd/system/docker-insight.service; enabled; vendor preset: disabled)
Active: inactive (dead) since <Date and Time since it was last active>
```

Restarting Cisco VIM Unified Management

Step 1 Log In to the server in which the Insight container was stopped.

Step 2 Restart the Insight service by running the following command from the shell window:

```
systemctl restart docker-insight
```

a) Check the status of Insight container by running the following command: **docker ps -a | grep insight**.

```
STATUS
Up 6 seconds
```

b) Check the status of the service by running the following command:

```
systemctl status docker-insight
```

The following output is displayed:

```
Docker-insight.service - Insight Docker Service
Loaded: loaded (/usr/lib/systemd/system/docker-insight.service; enabled; vendor preset: disabled)
Active: active (running) since <Date and Time when it got active.>
```

Restoring VIM Insight

Cisco VIM Unified Management can be restored to its previous running state which existed at the time of backup.



Note We recommend you not to run the Unified Management on the node on which restore operation is performed.

Step 1 Re-image the Insight management node with the ISO version with which you want to restore the node, and with the same IP address that is used before the failure of the node.

Note Skip Step 1 if re-image is done with the ISO during management node restore. Insight restore can also be performed without re-image with ISO. Uninstall the Insight through `bootstrap_insight.py` and then restoring it by following below mentioned steps but this needs to be only done when you face issues with Insight and not in case of management node failure.

Step 2 Navigate to `/var/cisco/insight_backup/` directory at the remote server where the backup directory was copied during the backup operation.

Step 3 Copy the backup file to the `/var/cisco/insight_backup/` directory of the re-imaged management node. For example, to copy the backup directory from the remote host 20.0.0.5 to the management node `/var/cisco/insight_backup/directory`, execute the following command sequence: `rsync -e ssh -rtvpX --numeric-ids root@20.0.0.5:/var/cisco/insight_backup/backup_2017-01-09_14-04-38 /var/cisco/insight_backup`.

Step 4 In `/var/cisco/insight_backup/backup_<date-time>` directory, execute the following command:

```

# ./insight_restore -h

insight_restore : Cisco VIM Insight Restore Script
-----

Usage: ./insight_restore

-v          : Enable verbose mode

-h          : To display this help message

# ./insight_restore
This will initiate an Insight install with the backed up data.

VIM Insight restore logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log

Management Node Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+

Insight standalone Input Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
+-----+-----+-----+

Setting up Insight, Kindly wait!!!
Cisco VIM Insight Installed successfully!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM UI Admin Email ID | PASS | Check for info @: <abs path of insight_setup_data.yaml> |
| | | |
| VIM UI Admin Password | PASS | Check for info @ /opt/cisco/insight/secrets.yaml |
| VIM Insight Workspace | PASS | /root/<insight_ws> |
+-----+-----+-----+

Cisco VIM Insight Autobackup Service Info!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+

```

```
| VIM Insight Autobackup | PASS | [ACTIVE]: Running 'insight-autobackup.service' |
+-----+-----+-----+
VIM Insight restore successfully completed!

Done with VIM Insight restore!
VIM Insight restore logs are at: /var/log/insight/bootstrap_insight/

As the summary table describes, your VIM Insight workspace is restored and hence you need to use
bootstrap_insight.py from the mentioned workspace for performing any actions from here on.
```

Step 5 Run the following command to verify Insight status after restore operation.

```
# cd /root/<insight_ws>
# ./bootstrap_insight.py -a install-status
                        Cisco VIM Insight Install Status!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight Version | PASS | <release_tag> |
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM Insight Container | PASS | insight_<tag_id> |
| VIM Mariadb Container | PASS | mariadb_<tag_id> |
| VIM Insight Autobackup | PASS | [ACTIVE]: Running 'insight-autobackup.service' |
| VIM Insight Workspace | PASS | /root/installer-<tag_id>/insight |
+-----+-----+-----+
```

Reconfiguring VIM Unified Management

UM reconfigure action provides you with three major functionalities:

1. Reconfigure Insight TLS Certificate.
2. Switch from Self Signed TLS Certificate to 3rdParty TLS Certificate.
3. Reconfigure Insight MySQL Database Password.



Note Unified Management reconfigure is not allowed after an update, as the update is an intermediate stage between rollback and commit.

Reconfiguring Insight TLS Certificate

As the Insight web-service is protected by TLS, hence reconfigure action provides flexibility to change the existing TLS Certificate. As there were two approaches to configure it, there are also two approaches to change it.

Reconfiguring 3rd Party TLS Certificate

If you had provided your own TLS Certificate before Insight Installation through PEM_PATH key in insight_setup_data.yaml, then perform the following steps to reconfigure it.

Step 1 # cd <path insight_setup_data.yaml>

Step 2 Open insight_setup_data.yaml file # vi insight_setup_data.yaml

Step 3 Edit the insight_setup_data.yaml to change value of PEM_PATH key to point to path of your new valid TLS Certificate. Save the file after editing. For example: PEM_PATH: "/root/new_tls.pem"

Note Only change in value of the PEM_PATH key is allowed during reconfigure. Any change that is detected in insight_setup_data.yaml other than PEM_PATH value results in failure of Insight reconfigure.

Step 4 # cd <insight_ws>

```
# ./bootstrap_insight.py -a reconfigure -f <path_to insight_setup_data.yaml>
```

VIM Insight reconfigure logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log

Perform the action. Continue (Y/N)y

Management Node Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+
```

Insight standalone Input Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+-----+-----+
```

WARNING!! reconfigure will have few secs of Outage for Insight!

Cisco VIM Insight Already Exists!

```
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM UI Admin Email ID | PASS | Check for info @: <abs path of insight_setup_data.yaml> |
| | | |
| VIM UI Admin Password | PASS | Check for info @ /opt/cisco/insight/secrets.yaml |
| VIM Insight Workspace | PASS | /root/<insight_ws> |
+-----+-----+-----+
```

```

Cisco VIM Insight backup Info!
+-----+-----+-----+
| Description          | Status | Details |
+-----+-----+-----+
| Insight backup Status | PASS   | Backup done @ |
|                       |        | /var/cisco/insight_backup/backup-<release_tag>-<date_time> |
+-----+-----+-----+

Done with VIM Insight reconfigure!
VIM Insight reconfigure logs are at: "/var/log/insight/bootstrap_insight/"

As the summary table describes Insight gets autobacked up after reconfigure at /var/cisco/insight_backup
to preserve the latest state of Insight.

```

Reconfiguring Self Signed TLS Certificate

If you had created a new TLS Certificate through `tls_insight_cert_gen.py` before Insight Installation then perform the following steps to reconfigure it.

Step 1 Run the following commands to reconfigure the self signed TLS certificate:

```

# cd <insight_ws>
# ./tls_insight_cert_gen.py -h
usage: tls_insight_cert_gen.py [-h] [--overwrite] --file INSIGHTSETUPDATA
TLS cert generator Insight

optional arguments:
  -h, --help            show this help message and exit
  --overwrite, -o      Overwrite Insight certificates if already present in  openstack config directory
  --file INSIGHTSETUPDATA, -f INSIGHTSETUPDATA
                        Location of insight_setup_data.yaml
                        # ./tls_insight_cert_gen.py -f <path insight_setup_data.yaml> --overwrite
                        This will overwrite the existing TLS certificate.

```

```

Management Node Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+

Insight standalone Input Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+

```

```

| Insight standalone Schema Validation      | PASS | None |
| Valid Key Check in Insight Setup Data    | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Test Email Server for Insight            | PASS | None |
+-----+-----+-----+

```

Generating a 4096 bit RSA private key

```

.....++
.....++
writing new private key to '../openstack-configs/insight.key'

```

Step 2 Following are the steps to run the bootstrap:

```

# ./bootstrap_insight.py -a reconfigure -f <path_to_insight_setup_data.yaml>
VIM Insight reconfigure logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log

```

Perform the action. Continue (Y/N)y

Management Node Validations!

```

+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+

```

Insight standalone Input Validations!

```

+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
| Test Email Server for Insight | PASS | None |
+-----+-----+-----+

```

WARNING!! reconfigure will have few secs of Outage for Insight
Cisco VIM Insight Already Exists!

```

+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM UI Admin Email ID | PASS | Check for info @: <abs path of insight_setup_data.yaml> |
| VIM UI Admin Password | PASS | Check for info @ /opt/cisco/insight/secrets.yaml |
| VIM Insight Workspace | PASS | /root/<insight_ws> |
+-----+-----+-----+

```

Cisco VIM Insight backup Info!

```

+-----+-----+-----+

```

Switch from Self Signed TLS Certificate to 3rd Party TLS Certificate

```

| Description          | Status | Details
|-----+-----+-----+
| Insight backup      | PASS   | Backup done @
|                     |        | /var/cisco/insight_backup/insight_backup_<release_tag>_<date_time>|
+-----+-----+-----+

```

Done with VIM Insight reconfigure!
 VIM Insight reconfigure logs are at: "/var/log/insight/bootstrap_insight/"

As the summary table describes Insight gets autobacked up after reconfigure at /var/cisco/insight_backup to preserve the latest state of Insight.

Switch from Self Signed TLS Certificate to 3rd Party TLS Certificate

If you had created a new TLS Certificate through `tls_insight_cert_gen.py` before Insight Installation and now want to switch to your own TLS Certificate then perform the following steps.



Note You cannot switch from 3rd Party TLS Certificate to Self Signed TLS Certificate.

Step 1 To switch from self-signed TLS Certificate to 3rd party TLS certificate open and edit the `insight_setup_data.yaml` to add `PEM_PATH` by running the following command:

```

# cd <path insight_setup_data.yaml>
# vi insight_setup_data.yaml
Edit the insight_setup_data.yaml to add PEM_PATH key to point to path of your new valid TLS Certificate.
Save the file after editing.
For example:
PEM_PATH: "/root/new_tls.pem"

```

Step 2 Following is the command to run the bootstrap:

```

# cd <insight_ws>
# ./bootstrap_insight.py -a reconfigure -f <path_to insight_setup_data.yaml>

VIM Insight reconfigure logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log
Perform the action. Continue (Y/N)y

```

```

Management Node Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |

```

```

| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+

```

Insight standalone Input Validations!

```

+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
| Test Email Server for Insight | PASS | None |
+-----+-----+-----+

```

WARNING!! reconfigure will have few secs of Outage for Insight!

Cisco VIM Insight Already Exists!

```

+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM UI Admin Email ID | PASS | Check for info @: <abs path of insight_setup_data.yaml> |
| VIM UI Admin Password | PASS | Check for info @ /opt/cisco/insight/secrets.yaml |
| VIM Insight Workspace | PASS | /root/<insight_ws> |
+-----+-----+-----+

```

Cisco VIM Insight backup Info!

```

+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| Insight backup Status | PASS | Backup done @ |
| | | /var/cisco/insight_backup/insight_backup_<release_tag>_<date_time> |
+-----+-----+-----+

```

Done with VIM Insight reconfigure!

VIM Insight reconfigure logs are at: "/var/log/insight/bootstrap_insight/"

As the summary table describes Insight gets autobacked up after reconfigure at /var/cisco/insight_backup to preserve the latest state of Insight.

Reconfiguring Insight MySQL Database Password

There are two approaches to reconfigure the MySQL DB password:

1. System generated Insight DB password.
2. User supplied Insight DB password.

System generated Insight DB password

Following are the steps to generate MySQL Insight DB password:

Step 1 To generate the Insight DB Password run the following command:

```
# cd <insight_ws>
# ./bootstrap_insight.py -a reconfigure -f <path_to insight_setup_data.yaml> --regenerate_secrets
```

VIM Insight reconfigure logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log
Perform the action. Continue (Y/N)y

Management Node Validations!

```
+-----+
| Rule | Status | Error |
+-----+
| Check Kernel Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+
```

Insight standalone Input Validations!

```
+-----+
| Rule | Status | Error |
+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+
```

WARNING!! reconfigure will have few secs of Outage for Insight!

Cisco VIM Insight Already Exists!

```
+-----+
| Description | Status | Details |
+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> | |
| VIM UI Admin Email ID | PASS | Check for info @: <abs path of insight_setup_data.yaml> |
| | | | |
| VIM UI Admin Password | PASS | Check for info @ /opt/cisco/insight/secrets.yaml |
| VIM Insight Workspace | PASS | /root/<insight_ws> |
+-----+
```

Cisco VIM Insight backup Info!

```
+-----+
| Description | Status | Details |
+-----+
| Insight backup Status | PASS | Backup done @ |
| | | /var/cisco/insight_backup/backup-<release_tag>-<date_time> |
+-----+
```

Done with VIM Insight reconfigure!

VIM Insight reconfigure logs are at: "/var/log/insight/bootstrap_insight/"

As the summary table describes Insight gets autobacked up after reconfigure at /var/cisco/insight_backup to preserve the latest state of Insight.

Step 2 Verify the password change by running the following command:

```
# cat /opt/cisco/insight/secrets.yaml
DB_ROOT_PASSWORD: <new_db_password>
```

User supplied Insight DB password

Step 1 To provide your own MYSQL DB Password follow the below steps:

Note Your new DB password must contain alphanumeric characters and should be at least 8 characters long.

```
# cd <insight_ws>
# ./bootstrap_insight.py -a reconfigure -f <path_to insight_setup_data.yaml> --setpassword

VIM Insight reconfigure logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log
Perform the action. Continue (Y/N)y
Password for DB_ROOT_PASSWORD: <enter_valid_db_password>
```

Management Node Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+
```

Insight standalone Input Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+-----+-----+
```

WARNING!! reconfigure will have few secs of Outage for Insight!

Cisco VIM Insight Already Exists!

```
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
```

```

| VIM Insight UI URL      | PASS | https://<br_api:9000> |
| VIM UI Admin Email ID | PASS | Check for info @: <abs path of insight_setup_data.yaml> |
| VIM UI Admin Password | PASS | Check for info @ /opt/cisco/insight/secrets.yaml |
| VIM Insight Workspace | PASS | /root/<insight_ws> |
+-----+-----+-----+

```

Cisco VIM Insight backup Info!

```

+-----+-----+-----+
| Description          | Status| Details |
+-----+-----+-----+
| Insight backup Status| PASS  | Backup done @
|                      |       | /var/cisco/insight_backup/insight_backup_<release_tag>_<date_time>|
+-----+-----+-----+

```

Done with VIM Insight reconfigure!

VIM Insight reconfigure logs are at: "/var/log/insight/bootstrap_insight/"

As the summary table describes Insight gets autobacked up after reconfigure at /var/cisco/insight_backup to preserve the latest state of Insight.

Step 2 Verify the password change by running the following command:

```

# cat /opt/cisco/insight/secrets.yaml
DB_ROOT_PASSWORD: <new_db_password>

```

Reconfiguring Unified Management SMTP Server

Unified Management requires a valid SMTP Server to send mails to users (Pod-Admin, UI-Admin, and regular users). If SMTP Server goes down, you can reconfigure it.

Following values can be reconfigured:

- INSIGHT_SMTP_SERVER
- INSIGHT_EMAIL_ALIAS_PASSWORD (only needed for Authenticated SMTP server)
- INSIGHT_EMAIL_ALIAS
- INSIGHT_SMTP_PORT (optional, defaults to 25)

Step 1 Following are the steps to reconfigure the SMTP server:

```

# cd <path insight_setup_data.yaml>
Open insight_setup_data.yaml file
# vi insight_setup_data.yaml
Edit the insight_setup_data.yaml to change value of INSIGHT_SMTP_SERVER key. Save the file after editing.

```

Step 2 Run the bootstrap command as follows:

```

# cd <insight_ws>
# ./bootstrap_insight.py -a reconfigure -f <path to insight_setup_data.yaml>
VIM Insight reconfigure logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log
Perform the action. Continue (Y/N)y

```

Management Node Validations!

Rule	Status	Error
Check Kernel Version	PASS	None
Check Ansible Version	PASS	None
Check Docker Version	PASS	None
Check Management Node Tag	PASS	None
Check Bond Intf. Settings	PASS	None
Root Password Check	PASS	None
Check Boot Partition Settings	PASS	None
Check LV Swap Settings	PASS	None
Check Docker Pool Settings	PASS	None
Check Home Dir Partition	PASS	None
Check Root Dir Partition	PASS	None
Check /var Partition	PASS	None
Check LVM partition	PASS	None
Check RHEL Pkgs Install State	PASS	None

Insight standalone Input Validations!

Rule	Status	Error
Insight standalone Schema Validation	PASS	None
Valid Key Check in Insight Setup Data	PASS	None
Duplicate Key Check In Insight Setup Data	PASS	None
Config Change Check in Insight Setup Data	PASS	None
Test Email Server for Insight	PASS	None

WARNING!! reconfigure will have few secs of Outage for Insight!

Cisco VIM Insight Already Exists!

Description	Status	Details
VIM Insight UI URL	PASS	https://<br_api:9000>
VIM UI Admin Email ID	PASS	Check for info @: <abs path of insight_setup_data.yaml>
VIM UI Admin Password	PASS	Check for info @ /opt/cisco/insight/secrets.yaml
VIM Insight Workspace	PASS	/root/<insight_ws>

Cisco VIM Insight backup Info!

Description	Status	Details
Insight backup Status	PASS	Backup done @
		/var/cisco/insight_backup/insight_backup_<release_tag>_<date_time>

Done with VIM Insight reconfigure!

VIM Insight reconfigure logs are at: "/var/log/insight/bootstrap_insight/"

As the summary table describes Insight gets autobacked up after reconfigure at /var/cisco/insight_backup to preserve the latest state of Insight.

Update VIM Insight

VIM Insight Update provides you the feature to switch to a new Insight release.

Following are some of the key points:

-
- Step 1** The update action will make the old docker containers of insight and mariadb in exit state and bring up new ones with the new tag.
 - Step 2** The old containers and images are restored until you perform the **Commit** action.
 - Step 3** Update is an intermediate action, from this state either you can do a **Commit** action to settle for the current version or do a **rollback** to revert back to the old version.
 - Step 4** The old workspace will be preserved in case you want to do a rollback to the previous version.
 - Step 5** After update your Insight workspace will be the new workspace you just extracted out of the tar ball.
 - Step 6** After an update operation, backup and reconfigure action is not allowed from either old or new Insight workspace.
- Note** Do not delete your old Insight workspace until you have successfully performed a **Commit** to the new Insight update.
-

Update Scenarios

Update action has the following scenarios:

- Insight and mariadb containers gets updated to a new tag.
- Either insight or mariadb container gets updated to a new tag.

Update VIM Insight with Internet Access

Following are the steps to update VIM Insight:

-
- Step 1** Get the new installer tar ball, which will be available after each release.
Extract the tar ball to get the new Insight workspace by running the following command:

```
# tar -xvzf mercury-installer.tar.gz
```

- Step 2** Update the VIM insight by running the following commands:

```
# cd /root/<new_insight_ws>/insight/  
/bootstrap_insight.py -a update
```

VIM Insight update logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log
Management Node Validations!

```
+-----+-----+-----+  
| Rule | Status | Error |  
+-----+-----+-----+  
| Check Kernel Version | PASS | None |  
| Check Docker Version | PASS | None |
```

```

| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+

Insight standalone Input Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+-----+-----+

Downloading Updated VIM Insight Artifacts, will take time!!!
Cisco VIM Insight update Info!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight Container: insight_<new_tag> | PASS | Updated from insight_<old_tag>|
| VIM Mariadb Container: mariadb_<new_tag> | PASS | Updated from mariadb_<old_tag>|
+-----+-----+-----+

Done with VIM Insight update!
VIM Insight update logs are at: "/var/log/insight/bootstrap_insight/"

```

Step 3 Verify the Insight Update.

```

# ./bootstrap_insight.py -a update-status
Cisco VIM Insight Update Status!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight Container: insight_<new_tag> | PASS | Updated from insight_<old_tag>|
| VIM Mariadb Container: insight_<new_tag> | PASS | Updated from mariadb_<old_tag>|
+-----+-----+-----+

```

VIM Insight without Internet Access

Step 1 Copy the new installer tar ball to the Insight Management Node.

Extract the tar ball to get the new Insight workspace by running the following command:

```
# tar -xvzf mercury-installer.tar.gz
```

Step 2 To download the new Insight artifacts follow the steps given in Cisco_NFVI_Install_Guide_2_2 under Preparing to Install Cisco NFVI on Management Nodes Without Internet Access, page 41.

Step 3 Run Import Artifacts:

```
# cd /root/installer_<tag_id>/tools
# ./import_artifacts.sh
```

This verifies that /var/cisco/artifacts on the management node has the following Insight artifacts, along with the other components 'insight-K9.tar' and 'mariadb-app-K9.tar'.'

Step 4 Update the Insight by running the following command:

```
# cd ../insight/
# ./bootstrap_insight.py -a update
```

VIM Insight update logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log
Management Node Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+
```

Insight standalone Input Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+-----+-----+
```

Updating VIM Insight, Kindly wait!!!

Cisco VIM Insight update Info!

```
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM Insight Container: insight_<new_tag> | PASS | Updated from insight_<old_tag> |
| VIM Mariadb Container: mariadb_<new_tag> | PASS | Updated from mariadb_<old_tag> |
| VIM Insight Workspace | PASS | /root/<new_insight_ws> |
+-----+-----+-----+
```

Done with VIM Insight update!

VIM Insight update logs are at: "/var/log/insight/bootstrap_insight/"

Step 5 Verify Insight Update by running the following command:

```
# ./bootstrap_insight.py -a update-status
Cisco VIM Insight Update Status!
```

```
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
```

```
| VIM Insight Container: insight_<new_tag> | PASS | Updated from insight_<old_tag> |
| VIM Mariadb Container: insight_<new_tag> | PASS | Updated from mariadb_<old_tag> |
+-----+-----+-----+
```

Rollback VIM Unified Management

VIM Unified Management Rollback provides feature to revert to the old UM release which is used before the update.

Following are some of the key points:

- The rollback action removes the new docker containers of insight and mariadb which is created after an update and bring up old ones with the old tag.
- The new workspace is used to update the operation later or the VIM may be running from it.
- After rollback, your Insight workspace is the old workspace which you were using before the update.

Following are the steps to perform Insight rollback:

Step 1 Run the following command to rollback VIM Insight:

```
# cd /root/<new_insight_ws>
# ./bootstrap_insight.py -a rollback
```

VIM Insight rollback logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log

Management Node Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+
```

Insight standalone Input Validations!

```
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+-----+-----+
```

```
VIM Insight rollback in progress, Kindly wait!!!
Cisco VIM Insight rollback Info!
+-----+
| Description                               | Status | Details                               |
+-----+-----+-----+
| VIM Insight UI URL                       | PASS   | https://<br_api:9000>                 |
| VIM Insight Container: insight_<old_tag> | PASS   | Rollback from insight_<new_tag>      |
| VIM Mariadb Container: mariadb_<old_tag> | PASS   | Rollback from mariadb_<new_tag>      |
| VIM Insight Workspace                    | PASS   | /root/<old_insight_ws>                |
+-----+-----+-----+
Done with VIM Insight rollback!
VIM Insight rollback logs are at: "/var/log/insight/bootstrap_insight/"
```

Step 2 Verify Rollback Status by running the following command:

```
# ./bootstrap_insight.py -a install-status
Cisco VIM Insight Install Status!
+-----+-----+-----+
| Description          | Status | Details                               |
+-----+-----+-----+
| VIM Insight Version | PASS   | <release_tag>                         |
| VIM Insight UI URL  | PASS   | https://<br_api:9000>                   |
| VIM Insight Container | PASS   | insight_<tag_id>                       |
| VIM Mariadb Container | PASS   | mariadb_<tag_id>                       |
| VIM Insight Workspace | PASS   | /root/<insight_ws>                     |
+-----+-----+-----+
```

Commit VIM Unified Management

VIM Insight Commit provides feature to settle for a new Insight release which you have been using after an update.

Following are some of the key points:

- The old workspace will not be deleted.
- After the commit, your Insight workspace is the new workspace which is used for the update.

Following are the steps to perform Insight commit:

Step 1 Run the following command to commit VIM Insight:

```
# cd /root/<new_insight_ws>
# ./bootstrap_insight.py -a commit
VIM Insight commit logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log Management Node
Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Check Kernel Version | PASS | None |
| Check Ansible Version | PASS | None |
| Check Docker Version | PASS | None |
| Check Management Node Tag | PASS | None |
```

```

| Check Bond Intf. Settings | PASS | None |
| Root Password Check | PASS | None |
| Check Boot Partition Settings | PASS | None |
| Check LV Swap Settings | PASS | None |
| Check Docker Pool Settings | PASS | None |
| Check Home Dir Partition | PASS | None |
| Check Root Dir Partition | PASS | None |
| Check /var Partition | PASS | None |
| Check LVM partition | PASS | None |
| Check RHEL Pkgs Install State | PASS | None |
+-----+-----+-----+

Insight standalone Input Validations!
+-----+-----+-----+
| Rule | Status | Error |
+-----+-----+-----+
| Insight standalone Schema Validation | PASS | None |
| Valid Key Check in Insight Setup Data | PASS | None |
| Duplicate Key Check In Insight Setup Data | PASS | None |
| Config Change Check in Insight Setup Data | PASS | None |
+-----+-----+-----+

VIM Insight commit in progress, Kindly wait!!!
Cisco VIM Insight commit Info!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM Insight Container: insight_<new_tag> | PASS | Old container: insight_<old_tag> removed |
| VIM Mariadb Container: mariadb_<new_tag> | PASS | Old container: mariadb_<old_tag> removed |
| VIM Insight Autobackup | PASS | [ACTIVE]: Running 'insight-autobackup.service'|
| VIM Insight Workspace | PASS | /root/<new_insight_ws> |
+-----+-----+-----+
Done with VIM Insight commit!
VIM Insight commit logs are at: "/var/log/insight/bootstrap_insight/"

```

Step 2 Verify Commit Status by running the following command:

```

# ./bootstrap_insight.py -a install-status Cisco VIM Insight Install Status!
+-----+-----+-----+
| Description | Status | Details |
+-----+-----+-----+
| VIM Insight Version | PASS | <release_tag> |
| VIM Insight UI URL | PASS | https://<br_api:9000> |
| VIM Insight Container | PASS | insight_<tag_id> |
| VIM Mariadb Container | PASS | mariadb_<tag_id> |
| VIM Insight Autoback | PASS |
| VIM Insight Workspace | PASS | [ACTIVE]: Running 'insight-autobackup.service'
| /root/<insight_ws> |
+-----+-----+-----+

```




CHAPTER 11

Overview to the Cisco Virtual Topology System

The Cisco Virtual Topology System (VTS) is an optional Cisco NFVI application that uses the Neutron driver and supports Cisco Vector Packet Processing. The following topics provide an overview to VTS architecture and features. When using VTS with Cisco NFVI, keep the following OpenStack tenant restrictions in mind:

Restriction	Description
Nova flavors: VM RAM > 512MB and equal to a multiple of 512MB	This limitation is due to NUMA and huge pages.
Nova Flavors: nova flavor-key m1.medium set hw:mem_page_size=large	VHOST mode is the only mode supported by the VTS installation at this time. To support VHOST connections nova needs the following configurations on each flavor that will be used.

- [Understanding Cisco VTS, on page 233](#)
- [Cisco VTS Architecture Overview, on page 234](#)
- [Virtual Topology Forwarder, on page 235](#)
- [Virtual Topology System High Availability, on page 237](#)

Understanding Cisco VTS

The Cisco Virtual Topology System (VTS) is a standards-based, open, overlay management and provisioning system for data center networks. It automates DC overlay fabric provisioning for both physical and virtual workloads.

Cisco VTS provides a network virtualization architecture and software-defined networking (SDN) framework that meets the requirements of multitenant data centers for cloud services. It enables a policy-based approach for overlay provisioning.

Cisco VTS automates complex network overlay provisioning and management tasks through integration with cloud orchestration systems such as OpenStack and VMware vCenter and abstracts out the complexity involved in managing heterogeneous network environments. The solution can be managed from the embedded Cisco VTS GUI or entirely by a set of northbound Representational State Transfer (REST) APIs that can be consumed by orchestration and cloud management systems.

Cisco VTS provides:

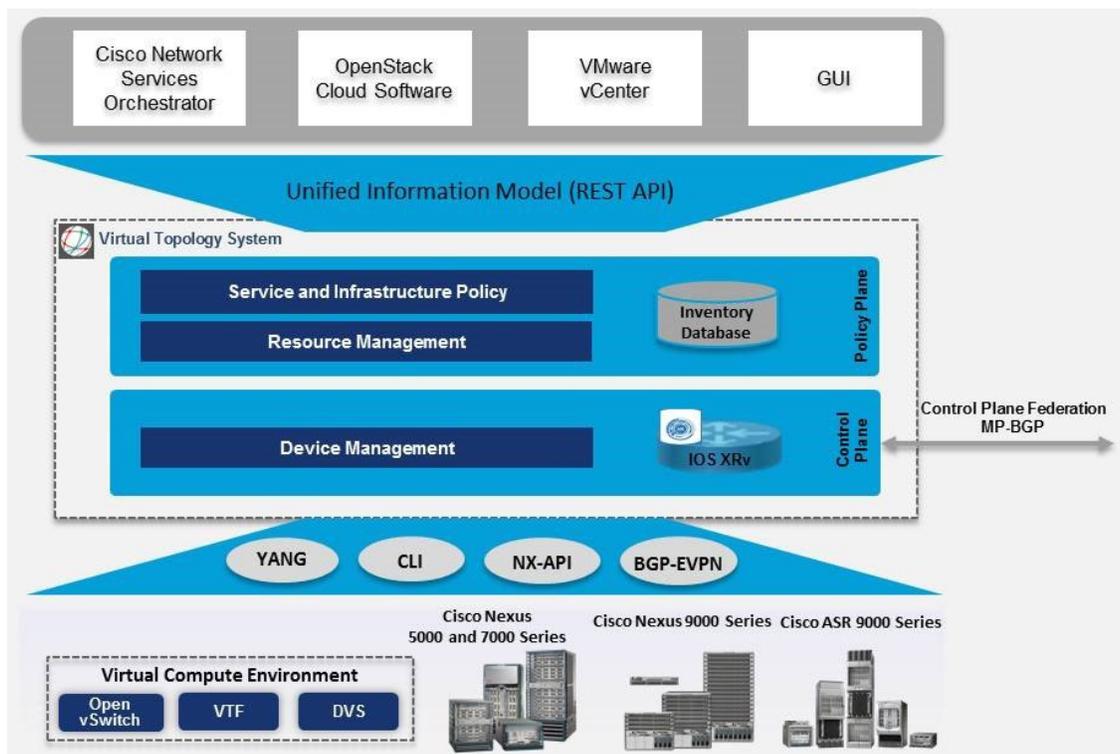
- Fabric automation

- Programmability
- Open, scalable, standards-based solution
- Cisco Nexus 2000, 3000, 5000, 7000, and 9000 Series Switches. For more information, see Supported Platforms in *Cisco VTS 2.6 Installation Guide*.
- Software forwarder (Virtual Topology Forwarder [VTF])

Cisco VTS Architecture Overview

Cisco VTS architecture has two main components: the Policy Plane and the Control Plane. These perform core functions such as SDN control, resource allocation, and core management function.

Figure 26: Cisco VTS Architecture



- **Policy Plane:** The policy plane enables Cisco VTS to implement a declarative policy model which is designed to capture intent and render of the user into a specific device-level construct. The solution exposes a set of modular policy constructs that can be flexibly organized into user-defined services for use cases across service provider and cloud environments. These policy constructs are exposed through a set of REST APIs that is consumed by orchestrators and applications to express user intent, or instantiated through the Cisco VTS GUI. Policy models are exposed as system policies or service policies.

System policies allow administrators to logically group devices into pods within or across data centers to define Admin Domains with common system parameters (for example, BGP-EVPN control plane with distributed Layer 2 and 3 gateways).

The inventory module maintains a database of the available physical entities (for example, data center interconnect [DCI] routers and top-of-rack leaf, spine, and border-leaf switches) and virtual entities (for example, VTFs) in the Virtual Topology System domain. The database also includes interconnections between these entities and details about all services instantiated within a Virtual Topology System domain.

The resource management module manages all available resource pools in the Virtual Topology System domain, including VLANs, VXLAN Network Identifiers (VNIs), IP addresses, and multicast groups.

- **Control Plane:** The control plane module serves as the SDN control subsystem that programs the various data planes including the VTFs residing on the x86 servers, hardware leafs, DCI gateways. The Control plane hosts Service Routing (SR) module, which provides routing services to Cisco VTS. The Service Routing (SR) module is responsible for calculating L2 and L3 tables and routes to provide connectivity between the different VMs for a given tenant and service chaining. The main components of this module are the VTSR and VTF. VTSR is the controller and Virtual topology forwarder (VTF) runs on each compute server hosting the tenant VMs.

Virtual Topology Forwarder

Virtual Topology Forwarder (VTF) runs on each compute server in the DC and provides connectivity to all tenant VMs hosted on the compute server. VTF supports both intra and inter DC/WAN connectivity. VTF allows Cisco VTS to terminate VXLAN tunnels on host servers by using the VTF as a Software VXLAN Tunnel Endpoint (VTEP). Cisco VTS also supports hybrid overlays by stitching together physical and virtual endpoints into a single VXLAN segment.

VTF has 2 major components—Cisco's VPP (Vector Packet Processing) and VPFA. VPFA is a Cisco agent running on each VMM compute resource. VPFA is the FIB agent which receives L2/L3 table forwarding information from VTSR to provide the connectivity to local tenant VMs hosted on its compute, and programs them in the VPP.

VTF is deployed as a virtual machine or in vhost mode, to deliver a high-performance software data plane on a host server.

Overview to Cisco VTF and VPP

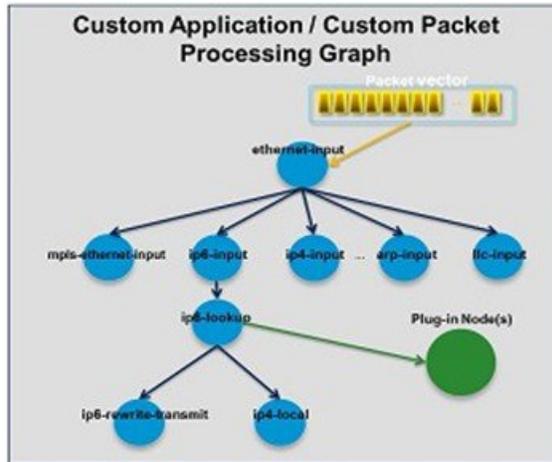
Cisco VTF is a Cisco Soft switch that is built on the Cisco Vector Packet Processing (VPP) technology.

The VPP platform is an extensible framework that provides productive and quality switch or router functionality. It is the open source version of the Cisco VPP technology, which is a high performance, packet-processing stack that can run on commodity CPUs.

The benefits of VPP are its high performance, proven technology, modularity, flexibility, and rich feature set.

The VPP platform is built on a packet-processing graph. This modular approach allows anyone to plugin new graph nodes. This makes extensibility rather simple, and the plugins can be customized for specific purposes.

Figure 27: VPP Platform



The VPP platform grabs all available packets from RX rings to form a vector of packets. A packet-processing graph is applied, node by node (including plugins) to the entire packet vector. Graph nodes are small and modular, and loosely coupled which makes it easy to include new graph nodes and rewire existing graph nodes.

A plugin can introduce new graph nodes or rearrange the packet-processing graph. You can also build a plugin independent from the VPP source and consider it as an independent component. A plugin can be installed by adding it to a plugin directory.

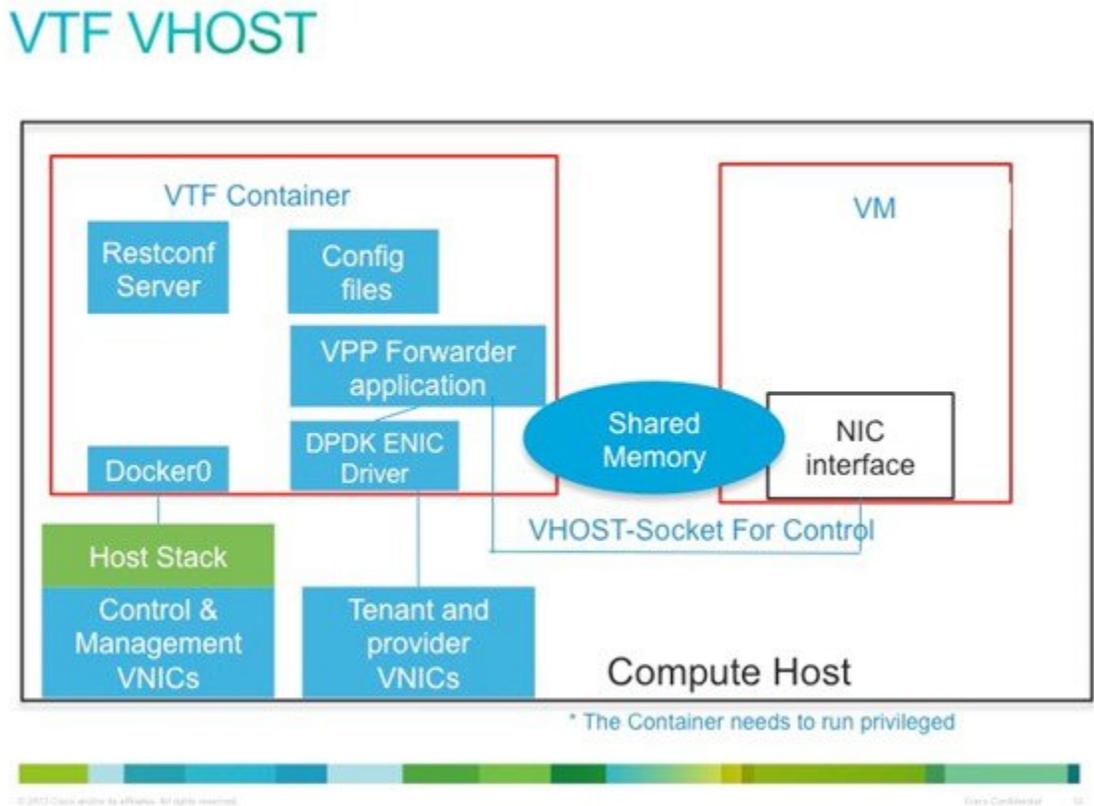
VTF uses remote plugin that binds into VPP using VPFA (VPF agent). The VPFA interacts with VPP application using low-level API. The VPFA exposes netconf or yang based API for remote devices to program the VTF through the VPFA.

VPP + VHOSTUSER

Vhost is a solution that allows the user space process to share a number of virtqueues directly with a Kernel driver. The transport mechanism in this case is the ability of the kernel side to access the user space application memory, and a number of ioeventfds and irqfds to serve as the kick mechanism. A QEMU guest uses an emulated PCI device, as the control plane to handle the QEMU. Once a virtqueue has been set up, the QEMU guest uses the Vhost API to pass direct control of a virtqueue to a Kernel driver.

In this model, a vhost_net driver directly passes the guest network traffic to a TUN device directly from the Kernel side, improving performance significantly.

Figure 28: VTF Vhost



In the above implementation, the guest NFV application directly writes packets into the TX rings, which are shared through a common vhost socket as the RX ring on the VPP. The VPP grabs these packets from the RX ring buffer and forwards the packets using the vector graphs it maintains.

Virtual Topology System High Availability

The Virtual Topology System solution is designed to support redundancy, with two solution instances running on separate hosts in an active-standby configuration.

During the initial setup, each instance is configured with both an underlay IP address and a virtual IP address. Virtual Router Redundancy Protocol (VRRP) is used between the instances to determine which instance is active.

The active-instance data is synchronized with the standby instance after each transaction to help ensure consistency of the control-plane information to accelerate failover after a failure. BGP peering is established from both Virtual Topology System instances for the distribution of tenant-specific routes. During the switchover, nonstop forwarding (NSF) and graceful restart help ensure that services are not disrupted.

See the *Installing VTS in High Availability Mode* section of the *Cisco VTS 2.6 Installation Guide* for the detailed procedure about setting up high availability.



CHAPTER 12

Managing Backup and Restore Operations

The following topics describe Cisco NFVI management node backup and restore operations.

- [Managing Backup and Restore Operations, on page 239](#)
- [Backing Up VIM UM, on page 241](#)
- [Restoring the Management Node, on page 245](#)
- [Management Node Autobackup, on page 247](#)

Managing Backup and Restore Operations

The management node hosts critical services such as Cisco VIM REST API, Cobbler for PXE, ELK for Logging/Kibana dashboard, and VMTP for the cloud validation in Cisco VIM.

The management node is not redundant during the initial Cisco VIM offering, hence it is recommended to take backup of the management node. Using the saved management node information, you can restore the management node if you are facing any issues with the platform.

Backing Up the Management Node

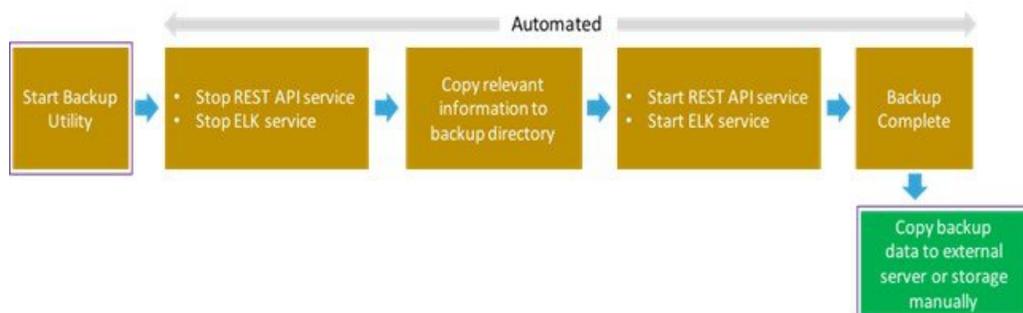
An administrator must maintain the number of backup snapshots on the management node. The backup of the management node is possible only after complete deployment of at least one Cisco VIM. Two copies of backup folders are maintained at the management node itself and the older copy will be overwritten when a next backup is performed.

During the backup operation, activities such as pod management, software update or upgrade, and addition or deletion or replacement of nodes cannot be performed.

The REST API and ELK services are stopped during the backup operation, the OpenStack Logs are cached on the control, compute, and storage nodes until the restoration of the management node is completed.

As part of the backup operation, two files are created: `.backup_files` and `.backup_hash`. `.backup_files` is a list of files that are backed up, while the second one is the hash. These two files are placed under the backup folder `/var/cisco/backup_<tag>_<date-time>` at the management node and also at the `/var/cisco/` folder of all three controllers. These two files are used during the restore validation. When you attempt to restore from a particular backup, these two files within this backup are compared to the files that are kept in the controllers. If there is any discrepancy, the restore validation fails and you are prompted to either terminate the restore operation or continue despite the validation failure. Only one copy of the `.backup_files` and `.backup_hash` are kept at the controllers, that is every time a new backup is created, these two files are overwritten with the most recent ones. Hence the restore validation passes only when the latest backup is used for restore.

Figure 29: Cisco NFVI Management Node Backup Operation

**Before you begin**

- Save the management node information (for example, IP address of the management node) for use during the restore operation.
- Ensure that you have the `br_mgmt` and `br_api` IP addresses and respective network information.

Step 1 Launch a SSH session to the Cisco NFVI management node.

Step 2 Navigate to the `<installer-ws>/tools/mgmt/` directory.

Step 3 Execute `mgmt_node_backup.py`.

What to do next

The backup operation takes approximately 30 minutes and creates the `backup_<tag>_<date-time>` directory in the `/var/cisco/` path.

Copy the directory to a remote server to recover the management node using `rsync`.

For example, to copy the backup directory to the remote server `20.0.0.5 /var/cisco/directory`, execute the following command sequence:

```
rsync -e ssh -go-rtvpX --numeric-ids/var/cisco/backup_2017-01-09_14-04-38
root@20.0.0.5:/var/cisco/
```



Note On the remote server, protect the backup directory for any unauthorized access as the backup files may contain sensitive information

At the remote server, change directory to where the backup directory is copied to; in this example `/var/cisco/backup_<tag>_<date-time>/`.

To verify if the backup is not corrupted or modified, execute `./check_integrity`.

`check_integrity` depends on the following packages, the packages are installed on the server where `check_integrity` is executed.

```
python-prettytable
python-jinja2
```

```
python-babel
python-markupsafe
python-setuptools
pytz
```

Backup with Forwarding ELK Logs to External Syslog Server

When the feature Forwarding ELK Logs to External Syslog Server is enabled, during the backup process, in both the autobackup and manual backup, the ELK Logs are not collected. For manual backups, you can override by appending the `-a` or `--add-elk` option to the backup command. The `-s` or `--skip-elk` option is to skip the ELK Logs collection regardless of the forwarding feature is enabled or not.

```
# cd installer/tools/mgmt
# ./mgmt_node_backup.py --help
Usage: ./mgmt_node_backup.py [options]
Options:
  -h, --help            show this help message and exit
  -s, --skip-elk        do not collect ELK logs during backup
  -a, --add-elk         force to also collect ELK logs on backup
```

Backing Up VIM UM

Administrator maintains the backup for Insight on the management node. The backup of the Insight is done only after the complete deployment of the Insight bootstrap. Only two copies of backup directory are maintained at the management node. The older copy is overwritten when a next Insight backup or autobackup takes place.

Insight backup is stored at the default backup location

`/var/cisco/insight_backup/insight_backup_<release_tag>_<date>_<time>`. If you want to take a backup of Insight at a different location use `-backupdir/-b` option from `bootstrap_insight`; details of which are provided later in this section.

Insight UI triggers an autobackup whenever it detects an operation relating to MySQL database entry to preserve the latest state of Insight.



Note Insight backup is not allowed after an update. Update is an intermediate stage between rollback and commit. Any change that is made relating to MySQL database entry after an update from UM UI is not backed up.

Autobackup Insight

If there is a change, Insight Installation automatically run a daemon process to take the autobackup.

Live status of the process is determined by checking the log located at `"/var/log/insight/insight_autobackup/insight_autobackup.logs"` or `systemctl status insight-autobackup`.



Note Max of 10-log files of size 1024*1024 are maintained in the directory.

Following are the scenarios where autobackup is initiated:

Insight Operation	Auto-backup Performed
Adding or Deleting POD	Yes
Changing POD REST Password and Certificate	Yes
Add/Edit/Delete all types of users	Yes
Add/Edit/Delete Roles	Yes
Modify User and Role association	Yes
Revoking or Adding user permission	Yes
Log in or Logout	No
Context Switching	No
Change User Password	Yes

Step 1 To check the status of the Insight perform the following steps:

```
systemctl status insight-autobackup
insight-autobackup.service - Insight Autobackup Service
  Loaded: loaded (/usr/lib/systemd/system/insight-autobackup.service; enabled; vendor preset:
disabled)
  Active: active (running) since Wed 2017-08-30 01:17:18 PDT; 19s ago
  Main PID: 19949 (python)
  Memory: 12.4M
  CGroup: /system.slice/insight-autobackup.service
          └─19949 /usr/bin/python /root/<installer-tag>/insight/playbooks/./insight_autobackup.py
```

Step 2 To stop Insight autobackup do the following:

```
systemctl stop insight-autobackup
insight-autobackup.service - Insight Autobackup Service
  Loaded: loaded (/usr/lib/systemd/system/insight-autobackup.service; enabled; vendor preset:
disabled)
  Active: inactive (dead) since Mon 2017-09-04 00:43:43 PDT; 5s ago
  Process: 19993 ExecStop=/bin/kill ${MAINPID} (code=exited, status=0/SUCCESS)
  Main PID: 19984
  Memory: 56.0K
  CGroup: /system.slice/insight-autobackup.service
```

Step 3 The following are the steps to start Insight autobackup:

```
systemctl start insight-autobackup
insight-autobackup.service - Insight Autobackup Service
  Loaded: loaded (/usr/lib/systemd/system/insight-autobackup.service; enabled; vendor preset:
disabled)
  Active: active (running) since Wed 2017-08-30 01:17:18 PDT; 19s ago
  Main PID: 19949 (python)
  Memory: 12.4M
  CGroup: /system.slice/insight-autobackup.service
          └─19949 /usr/bin/python /root/<installer-tag>/insight/playbooks/./insight_autobackup.py
```

Step 4 The way Insight works is as follows:

1. Install

- As soon as galera db and insight containers are up the script will be invoked.
- Log dir : tailf /var/log/insight/insight_autobackup_logs/insight_autobackup.log.
- It has a 10-seconds pulse which tells if the service is up or not.
 - [2017-09-04 00:49:01,504] INFO [Insight Autobackup] Insight Autobackup Service Running.
 - [2017-09-04 00:49:11,514] INFO [Insight Autobackup] Insight Autobackup Service Running.
 - [2017-09-04 00:49:21,525] INFO [Insight Autobackup] Insight Autobackup Service Running.
- If there is any change it takes a backup (time to check Sql diff is 30 seconds).
- It creates "rbac_latest.sql" and "insight_latest.tar.gz" and dump in the latest backup dir.
- During restore the bootstrap script checks if "rbac_latest.sql" or "insight_latest.tar.gz" is present in the backup dir.

2. Update

- During update bootstrap insight does not support backup.
- Autobackup service would be terminated and no backup would be maintained in the intermediate state.

3. Rollback

- Script are invoked again from the previous workspace.

4. Commit

- Script are invoked again from the new workspace.

5. Uninstall

- Service files are deleted.
- Log directory remains as the same.

Back Up Insight at Default Back Up Location

Step 1 Launch an SSH session to Cisco Insight management node and follow steps:

```
# cd <insight-ws>
#./bootstrap_insight.py -help

usage: bootstrap_insight.py [-h] --action ACTION
                          [--regenerate_secrets] [--setpassword]
                          [--file INSIGHTSETUPDATA] [--keep] [--verbose]
                          [--backupdir BACKUPDIR] [-y]

Insight install setup helper.
optional arguments:
  -h, --help            show this help message and exit
  --action ACTION, -a ACTION
                        install - Install Insight UI
```

```

install-status - Display Insight Install Status
reconfigure - Reconfigure Insight DB password or TLS Certificate
update - Update Insight UI
update-status - Display Insight Update Status
rollback - Rollback Insight UI update
commit - Commit Insight UI update
backup - Backup Insight UI
uninstall - Uninstall Insight UI

--regenerate_secrets, -r
System generated INSIGHT_DB_PASSWORD
--setpassword, -s User supplied INSIGHT_DB_PASSWORD,
--file INSIGHTSETUPDATA, -f INSIGHTSETUPDATA
Location of insight_setup_data.yaml
--keep, -k Preserve Insight artifacts during uninstall
--verbose, -v Verbose on/off
--backupdir BACKUPDIR, -b BACKUPDIR
Path to backup Insight
-y, --yes Option to skip reconfigure or uninstall steps without prompt

```

Step 2 Run the bootstrap command to view the Cisco VIM Insight backup details:

```

# ./bootstrap_insight.py -a backup
VIM Insight backup logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log

Cisco VIM Insight backup Info!
+-----+-----+-----+
| Description          | Status | Details |
|-----+-----+-----+
| Insight backup Status| PASS  | Backup done @
|-----+-----+-----+
|                       |        | /var/cisco/insight_backup/insight_backup_<release_tag>_<date_time>|
+-----+-----+-----+
Done with VIM Insight backup!

```

Backup Insight at user defined backup location

Step 1 Launch a SSH session to Cisco Insight management node and follow the below steps:

```

# cd <insight-ws>
# ./bootstrap_insight.py -help
usage: bootstrap_insight.py [-h] --action ACTION
                          [--regenerate_secrets] [--setpassword]
                          [--file INSIGHTSETUPDATA] [--keep] [--verbose]
                          [--backupdir BACKUPDIR] [-y]

Insight install setup helper.
optional arguments:
  -h, --help            show this help message and exit
  --action ACTION, -a ACTION
                        install - Install Insight UI
                        install-status - Display Insight Install Status
reconfigure - Reconfigure Insight DB password or TLS Certificate
update - Update Insight UI
update-status - Display Insight Update Status
rollback - Rollback Insight UI update
commit - Commit Insight UI update

```

```

        backup - Backup Insight UI
        uninstall - Uninstall Insight UI
--regenerate_secrets, -r
        System generated INSIGHT_DB_PASSWORD
--setpassword, -s
        User supplied INSIGHT_DB_PASSWORD,
--file INSIGHTSETUPDATA, -f INSIGHTSETUPDATA
        Location of insight_setup_data.yaml
--keep, -k
        Preserve Insight artifacts during uninstall
--verbose, -v
        Verbose on/off
--backupdir BACKUPDIR, -b BACKUPDIR
        Path to backup Insight
-y, --yes
        Option to skip reconfigure or uninstall steps without prompt

```

Step 2 Run the following command to view the Cisco VIM Insight backup details:

```

# ./bootstrap_insight.py -a backup --backupdir <user_defined_path>
VIM Insight backup logs are at: /var/log/insight/<bootstrap_insight_<date>_<time>.log

Cisco VIM Insight backup Info!
+-----+-----+-----+
| Description          | Status | Details          |
+-----+-----+-----+
| Insight backup Status | PASS   | Backup done @ <user_defined_path> |
+-----+-----+-----+
Done with VIM Insight backup!

```

What to do next

Copy the backup directory to a remote server using rsync to recover the Insight later. We recommend you to copy backup dir using rsync as it preserves the permissions of the files.

For example, to copy the backup directory to the remote server 20.0.0.5 /var/cisco/insight_backup/directory , execute the following command sequence: .

```

rsync -e ssh -rtvpX --numeric-ids
/var/cisco/insight_backup/insight_backup_2.1.5_2017-01-09_14-04-38
root@20.0.0.5:/var/cisco/insight_backup/

```

On the remote server, protect the backup directory for any unauthorized access, as the backup files may contain sensitive information

Restoring the Management Node

As an administrator, you have to reimage the management node with the same ISO version when the backup is performed, before initiating the restore operation. Restore fails when there is a version mismatch.

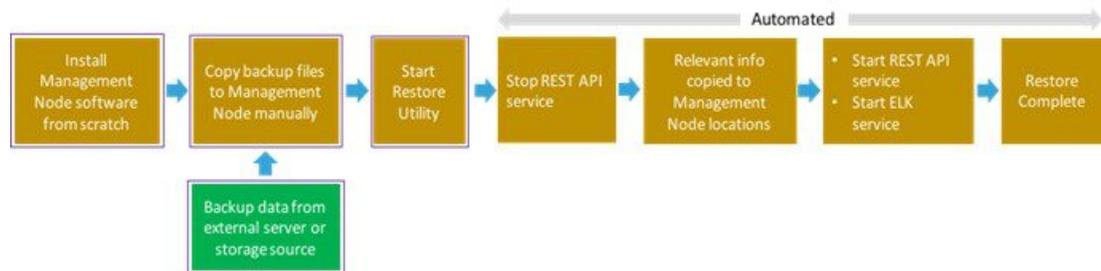


Note Version checking is available only for offline installation.

As part of the restore operation, system checks for the management node's IP address information to match the prior configuration. Logs are cached on the control, compute, and storage nodes from the moment of the management node fails until its restoration.

If you are using Cisco VIM Insight (in Tech Preview), in the same management node, you have to rebootstrap it for installation. During installation, RBAC and Pod registration information is lost, hence it is advised to make a note of the RBAC and Pod information.

Figure 30: Cisco NFVI Management Node Restore Operation



Before you begin

Ensure that you have the `br_mgmt` and `br_api` IP addresses of the failed management node.

Step 1 Reimage the management node with the ISO version with which you want to restore the node, and with the same IP address that is used before the failure of the node.

Step 2 Navigate to `/var/cisco/directory` at the remote server where the backup folder is copied during the backup operation. Execute `./check_integrity` to verify if the backup is not corrupted or modified.

Step 3 Copy the backup file to the `/var/cisco/directory` of the reimaged management node.

For example, to copy the backup folder from the remote host 20.0.0.5 to the management node `/var/cisco/directory`, execute the following command sequence:

```
rsync -e ssh -rtvpX --numeric-ids root@20.0.0.5:/var/cisco/backup_2017-01-09_14-04-38 /var/cisco/
```

Step 4 Navigate to the backup folder and execute the following command to verify if the backup is not corrupted or modified.

```
# cd /var/cisco/backup_<date-time>
# ./check-integrity
```

Step 5 In `/var/cisco/backup_<date-time>` folder, execute the following command:

```
/var/cisco/backup_<date-time> # ./restore
```

The restore operation takes around 45 minutes.

Step 6 Before restoration, the restore script performs validation of the backup folder. If validation fails, restore operation is halted and an error message is displayed. The script also verifies the last performed backup folder in the Management Node, and if any defects are detected, the you does confirm to proceed with restore operation.

```
...
2017-02-02 21:25:23 INFO Starting Cisco VIM restore...
2017-02-02 21:25:23 INFO Cisco VIM restore: estimated run time is approx. 45 mins...
2017-02-02 21:25:23 INFO Please see progress log for restore at
/var/log/mercury/installer/restore_2017-02-02_21:25:23.log
2017-02-02 21:25:27 ERROR Error: Backup id is not the one expected
Error: Found hashID file only in controller(s): j10-controller-2, j10-controller-3
Management backup files are ok (as per j10-controller-2)
Management backup files are ok (as per j10-controller-3)
The management node changed after the last backup was stored. Do you still want to proceed restoring
this management node? [Y/n] y
2017-02-02 22:17:55 INFO Workspace restored to /root/installer-6518
```

```
2017-02-02 22:17:55 INFO Cisco VIM restore: Executing restore playbook ...
2017-02-02 22:18:47 INFO Cisco VIM restore: Executing bootstrap playbook ...
```

Note The default behavior is to continue by keying **Return** or **Y**. Keying **N** terminates the restore operation.

```
...
2017-02-02 21:25:23 INFO Starting Cisco VIM restore...
2017-02-02 21:25:23 INFO Cisco VIM restore: estimated run time is approx. 45 mins...
2017-02-02 21:25:23 INFO Please see progress log for restore at
/var/log/mercury/installer/restore_2017-02-02_21:25:23.log
2017-02-02 21:25:27 ERROR Error: Backup id is not the one expected
Error: Found hashID file only in controller(s): j10-controller-2, j10-controller-3
Management backup files are ok (as per j10-controller-2)
Management backup files are ok (as per j10-controller-3)
The management node changed after the last backup was stored. Do you still want to proceed restoring
this management node? [Y/n] n
Aborting the restore operation as per user request
```

Once, restore operation ends, several health check points are automatically executed and the summary of results for that particular cloud availability is displayed.

Step 7 Run the following checks manually to verify the status of the restore:

- Check the status of the REST API server:

```
# cd installer-<tagid>/tools
# ./restapi.py -a status
Status of the REST API Server: active (running) since Thu 2016-08-18 09:15:39 UTC; 9h ago
REST API launch directory: /root/installer-<tagid>/
```

- Check the setup_data and runtime consistency of the management node:

```
# cd installer-<tagid>/; ciscovim run --perform 1,3 -y
```

- Execute the cloud sanity command:

```
# cd installer-<tagid>/tools
# ./cloud_sanity.py -c all
```

Management Node Autobackup

After the successful completion of certain Pod management operations, a backup of the management node is performed automatically. Only one copy of the autobackup folder is kept at /var/cisco/ at any given time. Directory format for the autobackup_<tag>_<timestamp>.

Following are the list of operations:

- Fresh install of Cisco VIM
- Commit an update
- Replace controller
- Add or Remove compute nodes
- Add or Remove the storage node
- Reconfigure

- NRVIMON

Enabling or disabling the variable autobackup, is defined in the setup_data.yaml file. It is enabled by default.

Add the following setup-data.yaml file snippet:

```
#####
# AutoBackup configuration
#####
#Default is True
#autobackup: True or False
```

The following tables shows when an auto-backup is performed during update or rollback or commit.

POD operation	Autobackup performed
Update	No
Rollback	No
Commit	Yes
Update fail with auto rollback	No

After creating a successful autobackup folder, you can copy it to an external server for later restoration as mentioned in [Restoring the Management Node, on page 245](#).

During the autobackup, if **Forwarding ELK Logs to the External Syslog server** option is enabled, the ElasticSearch database will not be maintained and the ELK Logs will not be recovered after restoring the management node.



CHAPTER 13

Troubleshooting

- [Displaying Cisco NFVI Node Names and IP Addresses, on page 249](#)
- [Verifying Cisco NFVI Node Interface Configurations, on page 250](#)
- [Displaying Cisco NFVI Node Network Configuration Files, on page 251](#)
- [Viewing Cisco NFVI Node Interface Bond Configuration Files, on page 252](#)
- [Viewing Cisco NFVI Node Route Information, on page 252](#)
- [Viewing Linux Network Namespace Route Information, on page 253](#)
- [Prior to Remove Storage Operation, on page 253](#)
- [Troubleshooting Cisco NFVI, on page 255](#)
- [Management Node Recovery Scenarios, on page 260](#)
- [Recovering Compute Node Scenario, on page 269](#)
- [Running the Cisco VIM Technical Support Tool, on page 271](#)
- [Tech-Support Configuration File, on page 272](#)
- [Tech-Support When Servers Are Offline, on page 274](#)
- [Disk-Maintenance Tool to Manage Physical Drives, on page 275](#)
- [OSD-Maintenance Tool, on page 278](#)
- [Utility to Resolve Cisco VIM Hardware Validation Failures, on page 281](#)
- [Cisco VIM Client Debug Option, on page 283](#)

Displaying Cisco NFVI Node Names and IP Addresses

Complete the following steps to display the Cisco NFVI node names and IP addresses.

Step 1 Log into the Cisco NFVI build node.

Step 2 The `openstack-configs/mercury_servers_info` file displays the node name and the address as follows.

```
# more openstack-configs/mercury_servers_info Total nodes: 5
Controller nodes: 3
+-----+-----+-----+-----+-----+-----+
| Server | CIMC | Management | Provision | Tenant | Storage |
+-----+-----+-----+-----+-----+-----+
| test-c-control-1 | 10.10.223.13 | 10.11.223.22 | 10.11.223.22 | 169.254.133.102 | None |
| | | | | |
| test-c-control-3 | 10.10.223.9 | 10.11.223.23 | 10.11.223.23 | 169.254.133.103 | None |
```

```

| | | | | | |
| test-c-control-2 | 10.10.223.10 | 10.11.223.24 | 10.11.223.24 | 169.254.133.104 | None |
| | | | | |
+-----+-----+-----+-----+-----+-----+
Compute nodes: 2
+-----+-----+-----+-----+-----+-----+
| Server | CIMC | Management | Provision | Tenant | Storage |
+-----+-----+-----+-----+-----+-----+
| test-c-compute-1 | 10.10.223.11 | 10.11.223.25 | 10.11.223.25 | 169.254.133.105 | None |
| | | | | |
| test-c-compute-2 | 10.10.223.12 | 10.11.223.26 | 10.11.223.26 | 169.254.133.106 | None |
| | | | | |
+

```

Note During the Cisco NFVI deployment, SSH public keys for each node are added to `.../.ssh/authorized_keys`, so you should be able to log in from the build node into each of the Cisco NFVI nodes without passwords. If, for some reason you do need account information, see the `openstack-configs/secrets.yaml` file on the build node.

Verifying Cisco NFVI Node Interface Configurations

Complete the following steps to verify the interface configurations of Cisco NFVI nodes:

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```

[root@mgmt-node~]# ssh root@control-server-1
[root@control-server-1 ~]#

```

Step 2 Enter the `ip a` command to get a list of all interfaces on the node:

```

[root@control-server-1 ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
2: enp8s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 54:a2:74:7d:42:1d brd ff:ff:ff:ff:ff:ff
3: enp9s0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 54:a2:74:7d:42:1e brd ff:ff:ff:ff:ff:ff
4: mx0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master mx state UP qlen 1000
    link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
5: mx1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master mx state UP qlen 1000
    link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
6: t0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master t state UP qlen 1000
    link/ether 54:a2:74:7d:42:23 brd ff:ff:ff:ff:ff:ff
7: t1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master t state UP qlen 1000
    link/ether 54:a2:74:7d:42:23 brd ff:ff:ff:ff:ff:ff
8: e0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master e state UP qlen 1000
    link/ether 54:a2:74:7d:42:25 brd ff:ff:ff:ff:ff:ff
9: e1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master e state UP qlen 1000
    link/ether 54:a2:74:7d:42:25 brd ff:ff:ff:ff:ff:ff
10: p0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master p state UP qlen 1000
    link/ether 54:a2:74:7d:42:27 brd ff:ff:ff:ff:ff:ff
11: p1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master p state UP qlen 1000
    link/ether 54:a2:74:7d:42:27 brd ff:ff:ff:ff:ff:ff
12: a0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master a state UP qlen 1000
    link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff

```

```

13: a1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master a state UP qlen 1000
    link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
14: bond0: <BROADCAST,MULTICAST,MASTER> mtu 1500 qdisc noop state DOWN
    link/ether 4a:2e:2a:9e:01:d1 brd ff:ff:ff:ff:ff:ff
15: a: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue master br_api state UP
    link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
16: br_api: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 54:a2:74:7d:42:29 brd ff:ff:ff:ff:ff:ff
17: e: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 54:a2:74:7d:42:25 brd ff:ff:ff:ff:ff:ff
18: mx: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue master br_mgmt state UP
    link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
19: br_mgmt: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 54:a2:74:7d:42:21 brd ff:ff:ff:ff:ff:ff
    inet 10.23.221.41/28 brd 10.23.221.47 scope global br_mgmt
        valid_lft forever preferred_lft forever
20: p: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 54:a2:74:7d:42:27 brd ff:ff:ff:ff:ff:ff
21: t: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 54:a2:74:7d:42:23 brd ff:ff:ff:ff:ff:ff
    inet 17.16.3.8/24 brd 17.16.3.255 scope global t
        valid_lft forever preferred_lft forever
22: docker0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state DOWN
    link/ether 02:42:70:f6:8b:da brd ff:ff:ff:ff:ff:ff
    inet 172.17.42.1/16 scope global docker0
        valid_lft forever preferred_lft forever
24: mgmt-out@if23: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br_mgmt state
    UP qlen 1000
    link/ether 5a:73:51:af:e5:e7 brd ff:ff:ff:ff:ff:ff link-netnsid 0
26: api-out@if25: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br_api state UP
    qlen 1000
    link/ether 6a:a6:fd:70:01:f9 brd ff:ff:ff:ff:ff:ff link-netnsid 0
    
```

Displaying Cisco NFVI Node Network Configuration Files

Complete the following steps to view a Cisco NFVI node network configuration files:

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```

[root@mgmt-node~]# ssh root@control-server-1
[root@control-server-1 ~]#
    
```

Step 2 List all of the network configuration files in the /etc/sysconfig/network-scripts directory, for example:

```

[root@control-server-1 ~]# ls /etc/sysconfig/network-scripts/
ifcfg-a          ifcfg-enp15s0  ifcfg-mx0      ifdown-ib      ifup           ifup-ppp
ifcfg-a0         ifcfg-enp16s0  ifcfg-mx1      ifdown-ipppp   ifup-aliases  ifup-routes
ifcfg-a1         ifcfg-enp17s0  ifcfg-p        ifdown-ipv6    ifup-bnep     ifup-sit
ifcfg-br_api     ifcfg-enp18s0  ifcfg-p0       ifdown-isdn    ifup-eth      ifup-Team
ifcfg-br_mgmt   ifcfg-enp19s0  ifcfg-p1       ifdown-post    ifup-ib       ifup-TeamPort
ifcfg-e          ifcfg-enp20s0  ifcfg-t        ifdown-ppp     ifup-ipppp    ifup-tunnel
ifcfg-e0         ifcfg-enp21s0  ifcfg-t0       ifdown-routes  ifup-ipv6     ifup-wireless
ifcfg-e1         ifcfg-enp8s0   ifcfg-t1       ifdown-sit     ifup-isdn     init.ipv6-global
ifcfg-enp12s0   ifcfg-enp9s0   ifdown        ifdown-Team    ifup-plip     network-functions
    
```

```
ifcfg-enp13s0  ifcfg-lo          ifdown-bnep  ifdown-TeamPort  ifup-plusb  network-functions-ipv6
ifcfg-enp14s0  ifcfg-mx          ifdown-eth   ifdown-tunnel    ifup-post
```

Viewing Cisco NFVI Node Interface Bond Configuration Files

Complete the following steps to view the Cisco NFVI node interface bond configuration files:

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```
[root@mgmt-node~]# ssh root@control-server-1
[root@control-server-1 ~]#
```

Step 2 List all of the network bond configuration files in the `/proc/net/bonding/` directory:

```
[root@control-server-1 ~]# ls /proc/net/bonding/
a  bond0  e  mx  p  t
```

Step 3 To view more information about a particular bond configuration, enter:

```
[root@control-server-1 ~]# more /proc/net/bonding/a
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: load balancing (xor)
Transmit Hash Policy: layer3+4 (1)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
```

```
Slave Interface: a0
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 1
Permanent HW addr: 54:a2:74:7d:42:29
Slave queue ID: 0
```

```
Slave Interface: a1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 2
Permanent HW addr: 54:a2:74:7d:42:2a
Slave queue ID: 0
```

Viewing Cisco NFVI Node Route Information

Complete the following steps to view Cisco NFVI node route information. Note that this is not the HAProxy container running on the controller. The default gateway should point to the gateway on the management network using the `br_mgmt` bridge.

Step 1 SSH into the target node, for example, one of the Cisco VIM controllers:

```
[root@mgmt-node~]# ssh root@control-server-1
[root@control-server-1 ~]#
```

Step 2 View the routing table (verify the default gateway) of the Cisco NFVI node:

```
[root@control-server-1 ~]# route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
0.0.0.0          10.23.221.33    0.0.0.0          UG    0      0      0 br_mgmt
10.23.221.32    0.0.0.0         255.255.255.240 U    0      0      0 br_mgmt
17.16.3.0       0.0.0.0         255.255.255.0   U    0      0      0 t
169.254.0.0     0.0.0.0         255.255.0.0     U    1016   0      0 br_api
169.254.0.0     0.0.0.0         255.255.0.0     U    1017   0      0 e
169.254.0.0     0.0.0.0         255.255.0.0     U    1019   0      0 br_mgmt
169.254.0.0     0.0.0.0         255.255.0.0     U    1020   0      0 p
169.254.0.0     0.0.0.0         255.255.0.0     U    1021   0      0 t
172.17.0.0      0.0.0.0         255.255.0.0     U    0      0      0 docker0
```

Viewing Linux Network Namespace Route Information

Complete the following steps to view the route information of the Linux network namespace that the HAProxy container uses on a Cisco NFVI controller node. The default gateway must point to the gateway on the API network using the API interface in the Linux network namespace.

Step 1 SSH into the target node. For example, one of the Cisco VIM controllers:

```
[root@mgmt-node~]# ssh root@control-server-1
[root@control-server-1 ~]#
```

Step 2 Enter the `ip netns` command to find the name of the network namespace:

```
[root@control-server-2 ~]# ip netns 17550 (id: 0)
```

Step 3 Enter the `ip netns exec` command to view the routing table (verify the default gateway) of the Linux network namespace:

```
[root@control-server-2 ~]# ip netns exec 17550 route -n
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref    Use Iface
0.0.0.0          172.29.86.1     0.0.0.0          UG    0      0      0 api
10.23.221.32    0.0.0.0         255.255.255.240 U    0      0      0 mgmt
172.29.86.0     0.0.0.0         255.255.255.0   U    0      0      0 api
```

Prior to Remove Storage Operation

Upon completion of the pod management operations such as add-storage, the operator has to ensure that any subsequent operation such as remove-storage on the same storage node is done after accounting for all of the

devices and their corresponding OSDs have been marked in the persistent crush map as shown in the output of the ceph osd crush tree.

Execute the following command on the storage node where a remove-storage pod operation is performed, to get a list of all the devices configured for ceph osds:

```
[root@storage-3 ~]$ df | grep -oh ceph-[0-9]*
[root@storage-3 ~]$ df | grep -oh ceph-[0-9]*
ceph-1
ceph-5
ceph-7
ceph-10
```

Login to any of the controller nodes and run the following commands within the ceph mon container:

```
$ cephmon
$ ceph osd crush tree
```

From the json output, locate the storage node to be removed and ensure all of the devices listed for ceph osds have corresponding osd entries for them by running the following commands:

```
{
  "id": -3,
  "name": "storage-3",
  "type": "host",
  "type_id": 1,
  "items": [
    {
      "id": 1,
      "name": "osd.1",
      "type": "osd",
      "type_id": 0,
      "crush_weight": 1.091095,
      "depth": 2
    },
    {
      "id": 5,
      "name": "osd.5",
      "type": "osd",
      "type_id": 0,
      "crush_weight": 1.091095,
      "depth": 2
    },
    {
      "id": 7,
      "name": "osd.7",
      "type": "osd",
      "type_id": 0,
      "crush_weight": 1.091095,
      "depth": 2
    },
    {
      "id": 10,
      "name": "osd.10",
      "type": "osd",
      "type_id": 0,
      "crush_weight": 1.091095,
      "depth": 2
    }
  ]
},
```

Troubleshooting Cisco NFVI

The following topics provide Cisco NFVI general troubleshooting procedures.

Managing CIMC and ISO Installation

When you are remote it is good to map the ISO through the CIMC Mapped vMedia.

To add new mapping:

Step 1 Click **Server > Remote Presence > Virtual Media > Add New Mapping**.

The screenshot shows a web-based dialog titled "Cisco IMC-Mapped vMedia" with a sub-header "Add New Map". The dialog contains the following fields and controls:

- Volume:** Text input field containing "f".
- Mount Type:** Dropdown menu set to "WWW(HTTP/HTTP)".
- Remote Share:** Text input field containing "http://172.29.172.152/downloads/".
- Remote File:** Text input field containing "buildnode-internal-13130.iso" and a "Browse..." button to the right.
- Mount Options:** Text input field containing "noauto".
- User Name:** Empty text input field.
- Password:** Empty text input field.

At the bottom right of the dialog are "Save" and "Close" buttons.

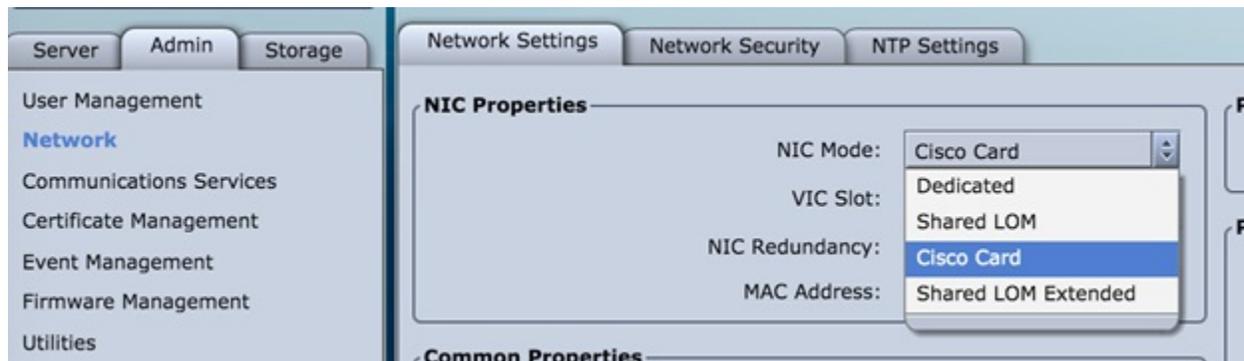
Step 2 Enter the field values such as the Volume, Mount Type, Remote Share, Remote File, User name, and Password.

Step 3 Click **Save**. The CIMC pulls the ISO directly from the HTTP server.

Management Node Installation Fails

Management node installation fails if the CIMC is configured for cisco card mode.

Choose the dedicated mode in the following screen:

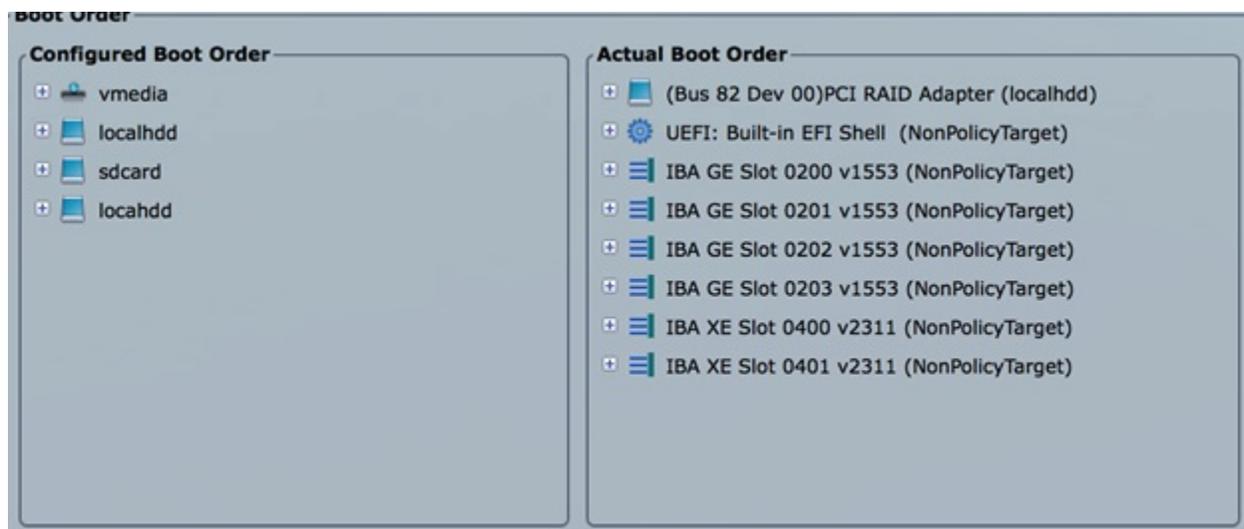


The selected method that is shown in the preceding screen is the incorrect mode.

Configuring Boot Order

Management node does not come up post reboot. It must boot from hard drive to check for the actual boot order.

Choose **Server > BIOS > Configure Boot Order > Boot Order**.

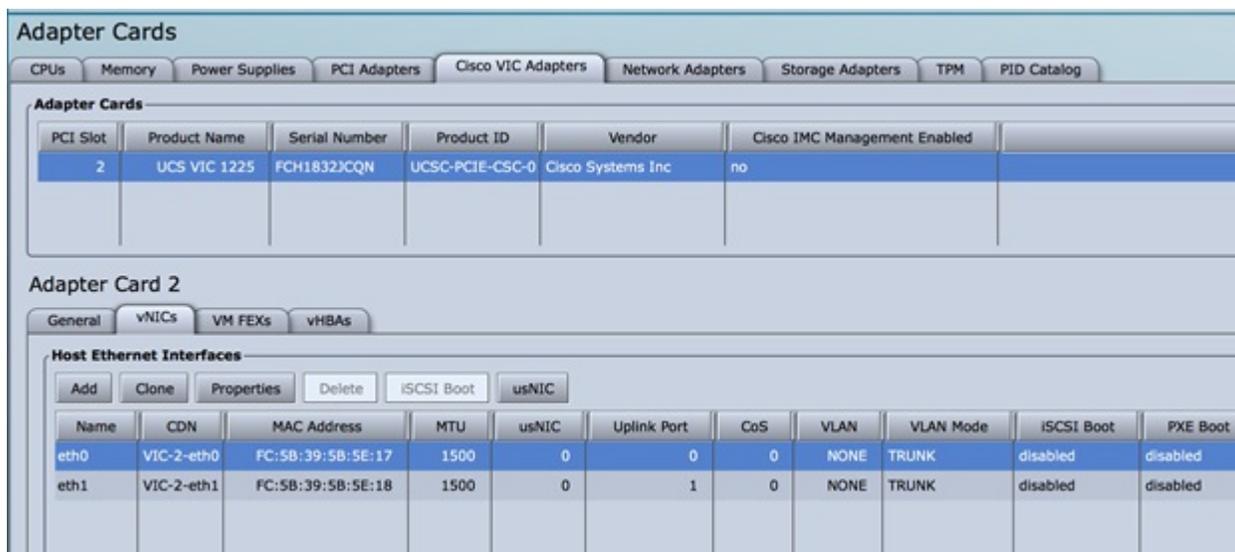


PXE Failure Issue During Baremetal Step

Perform the following steps in case of PXE boot failure:

- Step 1** Check log file `/var/log/mercury/mercury_baremetal_install.log` and connect to failing node CIMC KVM console to find out more on PXE boot failure reason.
- Step 2** Ensure all validations (step 1) and hardware validations (step 3) pass.
- Step 3** Check log file `/var/log/mercury/<UUID>/mercury_baremetal_install.log`.
- Step 4** Connect to KVM console of failing node(s) to find out more on PXE boot failure.
- Step 5** Check L2/L3 network connectivity between failing node(s) and management node.
- Step 6** Check for VPC configuration and port-channel status of failing node(s) and ensure `no lacp suspend-individual` is configured on the port-channel.
- Step 7** Check the actual PXE boot order must not differ from the boot-order configured.
- Step 8** Perform `tcpdump` on the management node interface `br_mgmt` to watch for UDP port 67 (dhcp) or UDP port 69 (tftp) `tcpdump -I br_mgmt port 67 or port 69 #` on the management node.
- Step 9** Perform `tcpdump` on the management node management interface `br_mgmt` on TCP 80 `tcpdump -I br_mgmt port 80 #` on the management node.
- Step 10** Check the apache log to watch the management IP address of failing node (if static allocated) `tail -f /var/log/cobblerhttpd/access_log #` on the management node.
- Step 11** For Authorization Required error messages during bare metal (Step 4) with CIMC operations such as hardware validations or cleaning up vNIC, check whether the maximum allowed simultaneous connection (4) are in use. All four connections are run when the 3rd party application monitoring CIMC does not properly close CIMC. This makes CiscoVIM installer not to log in using `xmlapi` with valid username and password. Check Cisco IMC logs on CIMC (Server > Faults and Logs > Cisco IMC Logs) for the reason why user was denied the access (maximum session, incorrect credentials.). The workaround is to disable 3rd party monitoring, wait at least 10 minutes and then perform CiscoVIM operations.
- Step 12** In case none of the nodes are getting DHCP address; DHCP requests arrive at the management node but no response goes out, then check CIMC VIC adapter settings. Server > Inventory > Cisco VIC Adapters > vNICs | VLAN & VLAN Mode. Ensure the VLAN (both id and mode) configured does not match with that of N9K switch

Option	Description
CIMC	Trunk:None
Switch	Access:vlan_mgmt



The following topics provide Cisco NFVI general troubleshooting procedures.

Container Download Problems

1. Check installer logs log file `/var/log/mercury/mercury_buildorchestration.log` for any build node orchestration failures including stuck "registry-Populate local registry". Downloading the Docker container from your management node can be slow.
2. Check the network connectivity between the management node and the remote registry in `defaults.yaml` on the management node (`grep "^registry:" openstack-configs/defaults.yaml`).
3. Verify valid remote registry credentials are defined in `setup_data.yaml` file.
4. A proxy server is required to pull the container images from remote registry. If a proxy is required, exclude all IP addresses for your setup including management node.

Cisco IMC Connection Problems during Bare Metal Installation

The cause may be Cisco IMC has too many connections, so the installer cannot connect to it. Clear the connections by logging into your Cisco IMC, going into the Admin->Sessions tab and clearing the connections.

API VIP Connection Problems

Verify the active HAProxy container is running in one of the controller nodes. On that controller within the HAProxy container namespace verify the IP address is assigned to the API interface. Also, verify that your ToR and the network infrastructure connecting your ToR is provisioned with API network segment VLAN.

HAProxy Services Downtime after Initial Installation or HA Failover

The HAProxy web interface can be accessed on TCP port 1936

```
http://<external_lb_vip_address>:1936/
Username: haproxy
Password: <HAPROXY_PASSWORD> from secrets.yaml file
```

After initial installation, the HAProxy web interface can report to several OpenStack services with downtime depending upon when that OpenStack service was installed after HAProxy install. The counters are not synchronized between HAProxy active and standby. After HA proxy failover, the downtime timers can change based on the uptime of new active HAProxy container.

Management Node Problems

Service Commands

To identify all the services that are running, enter:

```
$ systemctl -a | grep docker | grep service
  On controller ignore status of:
docker-neutronlb
  On compute ignore status of:
docker-neutronlb, docker-keystone
```

To start a service on a host, enter:

```
$ systemctl start <service_name>
```

To stop a service on a host, enter:

```
$ systemctl stop <service_name>
```

To restart a service on a host, enter:

```
$ systemctl restart <service_name>
```

To check service status on a host, enter:

```
$ systemctl status <service_name>
```

Connecting to Docker Container

To connect to the docket container do the following:

```
# generally, aliases are created for all containers
# use alias to identify those
alias | grep in_container
# checking specific alias by name
alias cobbler

# check docker containers
# alias created by CVIM
dp
# list docker containers
docker ps -a
# list docker images
docker images

# connecting to container
docker exec -it my_cobbler_<tag_id> /bin/bash

# connecting to docker container as privileged user
docker exec -it -u root my_cobbler_<tag_id> /bin/bash

# systemctl files
systemctl -a | egrep "docker-.*.service"

# check specific service
systemctl status mercury-restapi -l
systemctl status docker-vmtmp
```

```
# restart specific service
systemctl restart docker-vmtp
```

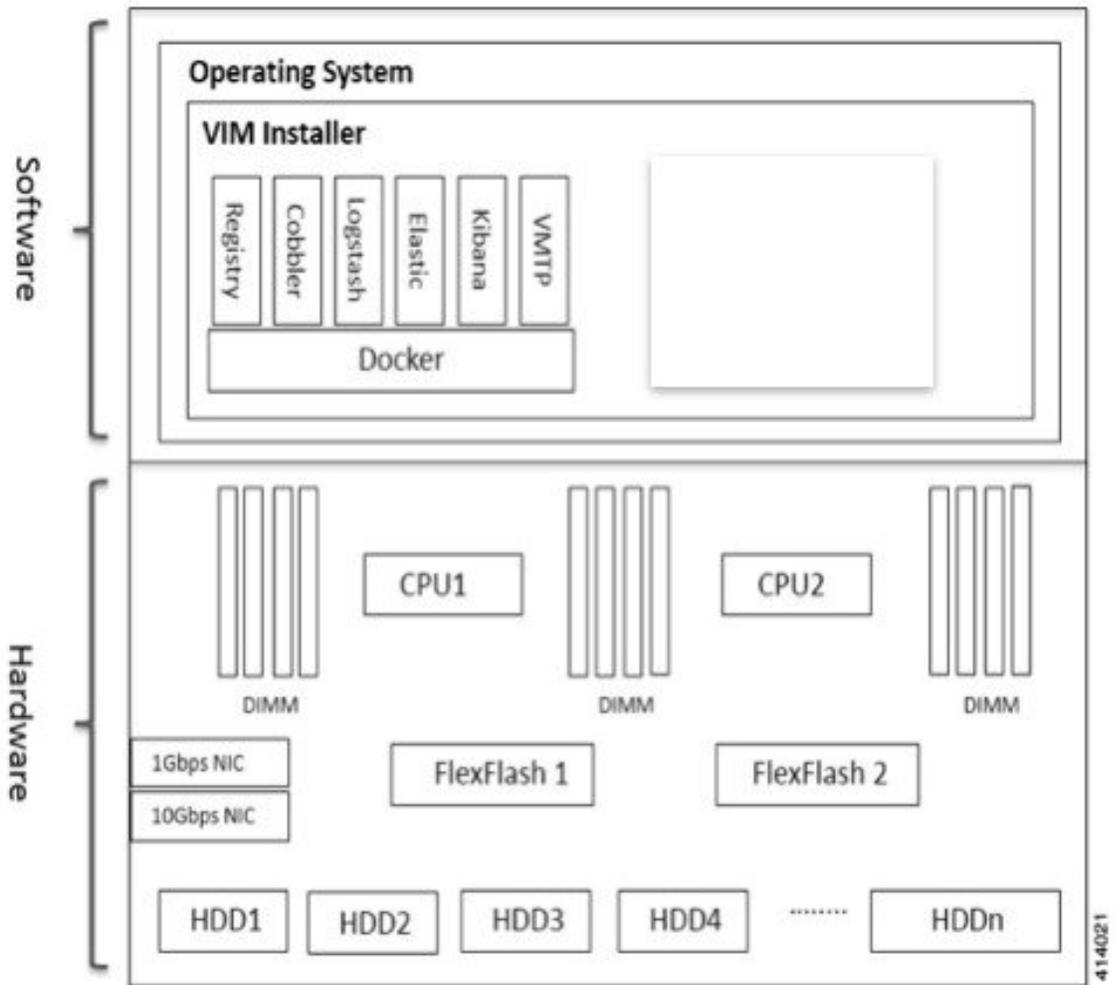
Management Node Recovery Scenarios

The Cisco NFVI management node hosts the Cisco VIM Rest API service, Cobbler for PXE services, ELK for Logging to Kibana dashboard services and VMTP for the cloud validation. As the maintenance node does not have redundancy, understanding its points of failure and recovery scenarios is important. Managing Node recovery scenarios are described in the following steps.

The management node architecture includes a Cisco UCS C240 M4 server with dual CPU socket. It has a 1-Gbps on-board (LOM) NIC and a 10-Gbps Cisco VIC mLOM. HDDs are used in 8,16, or 24 disk configurations.

The following figure shows the high-level maintenance node of the hardware and software architecture.

Figure 31: Cisco NFVI Management Node Architecture



Different management node hardware or software failures can cause Cisco NFVI service disruptions and outages. Some failed services can be recovered through manual intervention. In cases if the system is operational during a failure, double faults cannot be recoverable.

The following table lists the management node failure scenarios and their recovery options.

Table 4: Management Node Failure Scenarios

Scenario #	Failure or Trigger	Recoverable?	Operational Impact
1	Failure of 1 or 2 active HDD	Yes	No
2	Simultaneous failure of more than 2 active HDD	No	Yes
3	Spare HDD failure: 4 spare for 24 HDD; or 2 spare for 8 HDD	Yes	No

Scenario #	Failure or Trigger	Recoverable?	Operational Impact
4	Power outage/hard reboot	Yes	Yes
5	Graceful reboot	Yes	Yes
6	Docker daemon start failure	Yes	Yes
7	Service container (Cobbler, ELK) start failure	Yes	Yes
8	One link failure on bond interface	Yes	No
9	Two link failures on bond interface	Yes	Yes
10	REST API service failure	Yes	No
11	Graceful reboot with Cisco VIM Insight	Yes	Yes; CLI alternatives exist during reboot.
12	Power outage or hard reboot with Cisco VIM Insight	Yes	Yes
13	VIM Insight Container reinstallation	Yes	Yes; CLI alternatives exist during reinspection.
14	Cisco VIM Insight Container reboot	Yes	Yes; CLI alternatives exist during reboot.
15	Intel 1350 1Gbps LOM failure	Yes	Yes
16	Cisco VIC 1227 10-Gbps mLOM failure	Yes	Yes
17	DIMM memory failure	Yes	No
18	One CPU failure	Yes	No

Scenario 1: Failure of one or two active HDDs

The management node has either 8,16, or 24-HDDs. The HDDs are configured with RAID 6, which helps to enable data redundancy and storage performance and overcomes any unforeseen HDD failures.

- When 8 HDDs are installed, 7 are active disks and one is spare disk.
- When 16 HDDs are installed, 14 are active disks and two are spare disks.
- When 24 HDDs are installed, 20 are active disks and four are spare disks.

With RAID 6 up, two simultaneous active HDD failures can occur. When an HDD fails, the system begins automatic recovery by moving the spare disk to active state and begins recovering and rebuilding the new active HDD. It takes approximately 4 hours to rebuild the new disk and move to synchronized state. During this operation, the system is fully functional and no impacts are seen. However, you must monitor the system to ensure that more failures do not occur to enter into a double fault situation.

You can use the **storcli** commands to check the disk and RAID state as shown in the following commands:



Note Make sure that the node is running with hardware RAID by checking the storcli output and comparing to the one preceding.

```
[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show
<...snip...>
TOPOLOGY:
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID6 Optl N 4.087 TB dflt N N dflt N N
0 0 - - - RAID6 Optl N 4.087 TB dflt N N dflt N N <== RAID
6 in optimal state
0 0 0 252:1 1 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 1 252:2 2 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 2 252:3 3 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 3 252:4 4 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 4 252:5 5 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 5 252:6 6 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 6 252:7 7 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 - - 252:8 8 DRIVE DHS - 930.390 GB - - - - - N
-----
```

```
<...snip...>
PD LIST:
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:1 1 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U <== all disks
functioning
252:2 2 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
252:3 3 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:4 4 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:5 5 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:6 6 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:7 7 Onln 0 930.390 GB SAS HDD N N 512B ST91000640SS U
252:8 8 DHS 0 930.390 GB SAS HDD N N 512B ST91000640SS D
-----
```

```
[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show
<...snip...>
TOPOLOGY :
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID6 Pdgd N 4.087 TB dflt N N dflt N N <== RAID 6
in degraded state
0 0 - - - RAID6 Dgrd N 4.087 TB dflt N N dflt N N
0 0 0 252:8 8 DRIVE Rbld Y 930.390 GB dflt N N dflt - N
0 0 1 252:2 2 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 2 252:3 3 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 3 252:4 4 DRIVE Onln N 930.390 GB dflt N N dflt - N
```

```

0 0 4 252:5 5 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 5 252:6 6 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 6 252:7 7 DRIVE Onln N 930.390 GB dflt N N dflt - N
-----

```

<...snip...>

PD LIST :

=====

```

-----
EID:SlT DID State DG          Size Intf Med SED PI SeSz Model          Sp
-----
252:1      1 UGood - 837.258 GB SAS HDD N   N 512B ST900MM0006    U  <== active disk
in slot 1 disconnected from drive group 0
252:2      2 Onln  0 837.258 GB SAS HDD N   N 512B ST900MM0006    U
252:3      3 Onln  0 930.390 GB SAS HDD N   N 512B ST91000640SS   U
252:4      4 Onln  0 930.390 GB SAS HDD N   N 512B ST91000640SS   U
252:5      5 Onln  0 930.390 GB SAS HDD N   N 512B ST91000640SS   U
252:6      6 Onln  0 930.390 GB SAS HDD N   N 512B ST91000640SS   U
252:7      7 Onln  0 930.390 GB SAS HDD N   N 512B ST91000640SS   U
252:8      8 Rbld  0 930.390 GB SAS HDD N   N 512B ST91000640SS   U  <== spare disk
in slot 8 joined drive group 0 and in rebuilding state
-----

```

```

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s8 show rebuild
Controller = 0
Status = Success
Description = Show Drive Rebuild Status Succeeded.

```

```

-----
Drive-ID      Progress% Status          Estimated Time Left
-----
/c0/e252/s8      20 In progress 2 Hours 28 Minutes    <== spare disk in slot 8 rebuild
status
-----

```

To replace the failed disk and add it back as a spare:

```

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s1 add hotsparedrive dg=0
Controller = 0
Status = Success
Description = Add Hot Spare Succeeded.

```

```

[root@mgmt-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show

```

<...snip...>

TOPOLOGY :

=====

```

-----
DG Arr Row EID:Slot DID Type State BT          Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID6 Pdgd N 4.087 TB dflt N N dflt N N
0 0 - - - RAID6 Dgrd N 4.087 TB dflt N N dflt N N
0 0 0 252:8 8 DRIVE Rbld Y 930.390 GB dflt N N dflt - N
0 0 1 252:2 2 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 2 252:3 3 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 3 252:4 4 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 4 252:5 5 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 0 5 252:6 6 DRIVE Onln N 930.390 GB dflt N N dflt - N

```

```

0 0 6 252:7 7 DRIVE Onln N 930.390 GB dflt N N dflt - N
0 - - 252:1 1 DRIVE DHS - 837.258 GB - - - - N
-----

```

<...snip...>

PD LIST :
=====

```

-----
EID:SlT DID State DG          Size Intf Med SED PI SeSz Model          Sp
-----
252:1      1 DHS   0 837.258 GB SAS HDD N  N  512B ST900MM0006      U  <== replacement
  disk added back as spare
252:2      2 Onln  0 837.258 GB SAS HDD N  N  512B ST900MM0006      U
252:3      3 Onln  0 930.390 GB SAS HDD N  N  512B ST91000640SS      U
252:4      4 Onln  0 930.390 GB SAS HDD N  N  512B ST91000640SS      U
252:5      5 Onln  0 930.390 GB SAS HDD N  N  512B ST91000640SS      U
252:6      6 Onln  0 930.390 GB SAS HDD N  N  512B ST91000640SS      U
252:7      7 Onln  0 930.390 GB SAS HDD N  N  512B ST91000640SS      U
252:8      8 Rbld  0 930.390 GB SAS HDD N  N  512B ST91000640SS      U
-----

```

Scenario 2: Simultaneous failure of more than two active HDDs

If more than two HDD failures occur at the same time, the management node goes into an unrecoverable failure state because RAID 6 allows for recovery of up to two simultaneous HDD failures. To recover the management node, reinstall the operating system.

Scenario 3: Spare HDD failure

When the management node has 24 HDDs, four are designated as spares. Failure of any of the disks does not impact the RAID or system functionality. Cisco recommends replacing these disks when they fail (see the steps in Scenario 1) to serve as standby disks and so when an active disk fails, an auto-rebuild is triggered.

Scenario 4: Power outage or reboot

If a power outage or hard system reboot occurs, the system boots up, and come back to operational state. Services running on the management node during down time gets disrupted. See the steps in Scenario 9 for the list of commands to check the services status after recovery.

Scenario 5: System reboot

If a graceful system reboot occurs, the system boots up and come back to operational state. Services running on the management node during down time gets disrupted. See the steps in Scenario 9 for the list of commands to check the services status after recovery.

Scenario 6: Docker daemon start failure

The management node runs the services using Docker containers. If the Docker daemon fails to come up, it causes services such as ELK, Cobbler, and VMTP to go into down state. You can use the **systemctl** command to check the status of the Docker daemon, for example:

```

# systemctl status docker
docker.service - Docker Application Container Engine
Loaded: loaded (/usr/lib/systemd/system/docker.service; enabled; vendor preset: disabled)
Active: active (running) since Mon 2016-08-22 00:33:43 CEST; 21h ago
Docs: http://docs.docker.com
Main PID: 16728 (docker)

```

If the Docker daemon is in down state, use the **systemctl restart docker** command to restart the Docker service. Run the commands that are listed in Scenario 9 to verify that all the Docker services are active.

Scenario 7: Service container (Cobbler, ELK) start failure

As described in Scenario 8, all the services run as Docker containers on the management node. To find all services running as containers, use the **docker ps -a** command. If any services are in Exit state, use the **systemctl** command and **grep** for Docker to find the exact service name, for example:

```
# systemctl | grep docker- | awk '{print $1}'
docker-cobbler-tftp.service
docker-cobbler-web.service
docker-cobbler.service
docker-container-registry.service
docker-elasticsearch.service
docker-kibana.service
docker-logstash.service
docker-vmtp.service
```

If any services need restarting, use the **systemctl** command. For example, to restart a Kibana service:

```
# systemctl restart docker-kibana.service
```

Scenario 8: One link failure on the bond Interface

management node is set up with two different networks: `br_api` and `br_mgmt`. The `br_api` interface is the external. It is used for accessing outside services such as the Cisco VIM REST API, Kibana, and Cobbler. The `br_mgmt` interface is internal. It is used for provisioning and to provide management connectivity to all OpenStack nodes (control, compute and storage). Each network has two ports that are bonded to provide redundancy. If one port fails, the system remains completely functional through the other port. If a port fails, check for physical network connectivity, and remote switch configuration to debug the underlying cause of the link failure.

Scenario 9: Two link failures on the bond Interface

As described in Scenario 10, each network is configured with two ports. If both ports are down, the system is not reachable and management node services could be disrupted. After the ports are back up, the system is fully operational. Check the physical network connectivity and the remote switch configuration to debug the underlying link failure cause.

Scenario 10: REST API service failure

The management node runs the REST API service for Cisco VIM clients to reach the server. If the REST service is down, Cisco VIM clients cannot reach the server to trigger any server operations. However, with the exception of the REST service, other management node services remain operational.

To verify the management node REST services are fully operational, use the following command to check that the `httpd` and `mercury-restapi` services are in active and running state:

```
# systemctl status httpd
httpd.service - The Apache HTTP Server
  Loaded: loaded (/usr/lib/systemd/system/httpd.service; enabled; vendor preset: disabled)
  Active: active (running) since Mon 2016-08-22 00:22:10 CEST; 22h ago

# systemctl status mercury-restapi.service
mercury-restapi.service - Mercury Restapi
  Loaded: loaded (/usr/lib/systemd/system/mercury-restapi.service; enabled; vendor preset: disabled)
  Active: active (running) since Mon 2016-08-22 00:20:18 CEST; 22h ago
```

A tool is also provided so that you can check the REST API server status and the location of the folder it is running from. To execute run the following command:

```
# cd installer-<tagid>/tools
# ./restapi.py -a status
Status of the REST API Server: active (running) since Thu 2016-08-18 09:15:39 UTC; 9h
```

```
ago
  REST API launch directory: /root/installer-<tagid>/
```

Confirm the server status is active and check that the restapi launch folder matches the folder where the installation was launched. The restapi tool also provides the options to launch, tear down, and reset password for the restapi server as shown in the following command:

```
# ./restapi.py -h

usage: restapi.py [-h] --action ACTION [--yes] [--verbose]

REST API setup helper

optional arguments:
  -h, --help            show this help message and exit
  --action ACTION, -a ACTION
                        setup - Install and Start the REST API server.
                        teardown - Stop and Uninstall the REST API
                        server.
                        restart - Restart the REST API server.
                        regenerate-password - Regenerate the password for
                        REST API server.
                        reset-password - Reset the REST API password with
                        user given password.
                        status - Check the status of the REST API server
  --yes, -y            Skip the dialog. Yes to the action.
  --verbose, -v        Perform the action in verbose mode.
```

If the REST API server is not running, execute **ciscovim** to show the following error message:

```
# cd installer-<tagid>/
# ciscovim -setupfile ~/Save/<setup_data.yaml> run
```

If the installer directory or the REST API state is not correct or points to an incorrect REST API launch directory, go to the installer-<tagid>/tools directory and execute:

```
# ./restapi.py -action setup
```

To confirm that the REST API server state and launch directory is correct run the following command:

```
# ./restapi.py -action status
```

Scenario 11: Graceful reboot with Cisco VIM Insight

Cisco VIM Insight runs as a container on the management node. After a graceful reboot of the management node, the VIM Insight and its associated database containers comes up. So there is no impact on recovery.

Scenario 12: Power outage or hard reboot with VIM Insight

The Cisco VIM Insight container comes up automatically following a power outage or hard reset of the management node.

Scenario 13: Cisco VIM Insight reinstallation

If the management node which is running the Cisco VIM Insight fails and cannot come up, you must uninstall and reinstall the Cisco VIM UM. After the VM Insight container comes up, add the relevant bootstrap steps as listed in the install guide to register the pod. VIM Insight then automatically detects the installer status and reflects the present status appropriately.

To clean up and reinstall Cisco VIM UM run the following command:

```
# cd /root/installer-<tagid>/insight/
# ./bootstrap_insight.py -a uninstall -o standalone -f </root/insight_setup_data.yaml>
```

Scenario 14: VIM Insight Container reboot

On Reboot of the VIM Insight container, services continue to work as it is.

Scenario 15: Intel (I350) 1Gbps LOM failure

The management node is set up with an Intel (I350) 1-Gbps LOM for API connectivity. Two 1-Gbps ports are bonded to provide connectivity redundancy. No operational impact occurs if one of these ports goes down. However, if both ports fail, or the LOM network adapter fails, the system cannot be reached through the API IP address. If this occurs you must replace the server because the LOM is connected to the system motherboard. To recover the management node with a new server, complete the following steps. Make sure the new management node hardware profile, matches the existing server and the Cisco IMC IP address is assigned.

1. Shut down the existing management node.
2. Unplug the power from the existing and new management nodes.
3. Remove all HDDs from existing management node and install them in the same slots of the new management node.
4. Plug in the power to the new management node, but do not boot the node.
5. Verify the configured boot order is set to boot from local HDD.
6. Verify the Cisco NFVI management VLAN is configured on the Cisco VIC interfaces.
7. Boot the management node for the operating system to begin.

After the management node is up, the management node bond interface is down due to the incorrect MAC address. It points to old node network card MAC address.

8. Update the MAC address under `/etc/sysconfig/network-scripts`.
9. Reboot the management node.
It is fully operational. All interfaces has to be in an up state and be reachable.
10. Verify that Kibana and Cobbler dashboards are accessible.
11. Verify the Rest API services are up. See Scenario 15 for any recovery steps.

Scenario 16: Cisco VIC 1227 10Gbps mLOM failure

The management node is configured with a Cisco VIC 1227 dual port 10-Gbps mLOM adapter for connectivity to the other Cisco NFVI nodes. Two 10 Gbps ports are bonded to provide connectivity redundancy. If one of the 10-Gbps ports goes down, no operational impact occurs. However, if both Cisco VIC 10 Gbps ports fail, the system goes into an unreachable state on the management network. If this occurs, you must replace the VIC network adapters. Otherwise pod management and the Fluentd forwarding service is disrupted. If you replace a Cisco VIC, update the management and provisioning VLAN for the VIC interfaces using Cisco IMC and update the MAC address in the interfaces under `/etc/sysconfig/network-scripts` interface configuration file.

Scenario 17: DIMM memory failure

The management node is set up with multiple DIMM memory across different slots. Failure of one or memory modules could cause the system to go into unstable state, depending on how many DIMM memory failures occur. DIMM memory failures are standard system failures like any other Linux system server. If a DIMM memory fails, replace the memory module(s) as soon as possible to keep the system in stable state.

Scenario 18: One CPU failure

Cisco NFVI management nodes have dual core Intel CPUs (CPU1 and CPU2). If one CPU fails, the system remains operational. However, always replace failed CPU modules immediately. CPU failures are standard

system failures such as any other Linux system server. If a CPU fails, replace it immediately to keep the system in stable state.

Recovering Compute Node Scenario

The Cisco NFVI Compute node hosts the OpenStack services to provide processing, network, and storage resources to run instances. The node architecture includes a Cisco UCS C220 M4 server with dual CPU socket, 10-Gbps Cisco VIC mLOM, and two HDDs in RAID 1 configuration.

Failure of one active HDD

With RAID 1, data are shown and allows up to one active HDD failure. When an HDD fails, the node is still functional with no impacts. However, the data are no longer illustrated and losing another HDD results in unrecoverable and operational downtime. The failed disk has to be replaced soon as it takes approximately 2 hours to rebuild the new disk and move to synchronized state.

To check the disk and RAID state, run the storcli commands as follows:



Note Make sure that the node is running with hardware RAID by checking the storcli output and comparing to the one that is shown in the following command.

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show

<...snip...>

TOPOLOGY :
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID1 Opt1 N 837.258 GB dflt N N dflt N N <== RAID 1 in
optimal state
0 0 - - - RAID1 Opt1 N 837.258 GB dflt N N dflt N N
0 0 0 252:2 9 DRIVE Onln N 837.258 GB dflt N N dflt - N
0 0 1 252:3 11 DRIVE Onln N 837.258 GB dflt N N dflt - N
-----

<...snip...>

Physical Drives = 2

PD LIST :
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:2 9 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U <== all disks
functioning
252:3 11 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
-----

[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show

<...snip...>

TOPOLOGY :
```

```
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N <== RAID 1 in
degraded state.
0 0 - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N
0 0 0 - - DRIVE Msng - 837.258 GB - - - - - N
0 0 1 252:3 11 DRIVE Onln N 837.258 GB dflt N N dflt - N
-----
```

<...snip...>

PD LIST :

```
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:2 9 UGood - 837.258 GB SAS HDD N N 512B ST900MM0006 U <== active disk
in slot 2 disconnected from drive group 0
252:3 11 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
-----
```

To replace the failed disk and add it back as a spare run the following command:

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s2 add hotsparedrive dg=0
Controller = 0
Status = Success
Description = Add Hot Spare Succeeded.
```

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0 show
```

<...snip...>

TOPOLOGY :

```
=====
-----
DG Arr Row EID:Slot DID Type State BT Size PDC PI SED DS3 FSpace TR
-----
0 - - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N
0 0 - - - RAID1 Dgrd N 837.258 GB dflt N N dflt N N
0 0 0 252:2 9 DRIVE Rbld Y 837.258 GB dflt N N dflt - N
0 0 1 252:3 11 DRIVE Onln N 837.258 GB dflt N N dflt - N
-----
```

<...snip...>

PD LIST :

```
=====
-----
EID:SlT DID State DG Size Intf Med SED PI SeSz Model Sp
-----
252:2 9 Rbld 0 837.258 GB SAS HDD N N 512B ST900MM0006 U <== replacement
disk in slot 2 joined device group 0 and in rebuilding state
252:3 11 Onln 0 837.258 GB SAS HDD N N 512B ST900MM0006 U
-----
```

```
[root@compute-node ~]# /opt/MegaRAID/storcli/storcli64 /c0/e252/s2 show rebuild
Controller = 0
Status = Success
Description = Show Drive Rebuild Status Succeeded.
```

```

-----
Drive-ID      Progress% Status      Estimated Time Left
-----
/c0/e252/s2   10 In progress 1 Hours 9 Minutes  <== replacement disk in slot 2 rebuild
status
-----

```

Running the Cisco VIM Technical Support Tool

Cisco VIM includes a tech-support tool that you can use to gather Cisco VIM information to help solve issues working with Cisco Technical Support. The tech-support tool can be extended to execute custom scripts. It can be called after runner is executed at least once. The tech-support tool uses a configuration file that specifies what information to collect. The configuration file is located in the following location:

/root/openstack-configs/tech-support/tech_support_cfg.yaml.

The tech-support tool checks the point where the Cisco VIM installer has executed and collects the output of files or commands that is indicated by the configuration file. For example, if the installer fails at Step 3 (VALIDATION), the tech-support provides information that is listed in the configuration file up to Step 3 (included). You can override this default behavior by adding the --stage option to the command.

The tech-support script is located at the management node /root/installer-{tag-id}/tech-support directory. To run it after the runner execution, enter the following command:

```
./tech-support/tech_support.py
```

The command creates a compressed tar file containing all the information that is gathered. The file location is displayed in the console at the end of the execution. You need not have to execute the command with any options. However, if you want to override any default behavior, you can use the following options:

```

/tech_support.py --help
Usage: tech_support.py [options]

```

Tech-support collects information about your cloud

Options:

```

-h, --help          show this help message and exit
--stage=STAGE       specify the stage where installer left off
--config-file=CFG_FILE
                    specify alternate configuration file name
--tmp-dir=TMP_DIR   specify alternate temporary directory name
--file-size=TAIL_SIZE
                    specify max size (in KB) of each file collected
--host-list=HOST_LIST
                    List (comma separated) of the hostnames of the servers
                    to collect info from
--ip-list=IP_LIST   List (comma separated) of the IPv4 of the hosts to
                    collect info from
--exclude-mgmt-node
                    specify if mgmt node info needs to be excluded

```

Where:

- **stage**—tells at which state the installer left off. The possible values are: INPUT_VALIDATION, BUILDNODE_ORCHESTRATION, VALIDATION, BAREMETAL_INSTALL, COMMON_SETUP, CEPH, ORCHESTRATION or VMTP

- **config-file**—Provides the path for a specific configuration file. Make sure that your syntax is correct. Look at the default `/root/tech-support/openstack-configs/tech_support_cfg.yaml` file as an example on how to create a new `config-file` or modify the default file.
- **tmp-dir**—Provides the path to a temp directory `tech-support` can use to create the compressed tar file. The `tech-support` tool provides the infrastructure to execute standard Linux commands from packages that are included in the Cisco VIM installation. This infrastructure is extensible and you can add commands, files, or custom bash or Python scripts into the configuration file pane for the tool to collect the output of those commands or scripts. (See the README pane for more details.)
- **file-size**—Is an integer that specifies (in KB) the maximum file size that `tech-support` captures and tail the file if needed. By default, this value is set to 10 MB. For example, if no `file-size` option is provided and the `tech-support` has to collect `/var/log/mercury/data.log` and the `data.log` is more than 10 MB, `tech-support` gets the last 10 MB from `/var/log/mercury/data.log`.
- **host-list**: Provides the list of hosts one wants to collect from the `tech-support` through hostname defaults, to all hosts.
- **ip-list**: Provides the list of hosts one wants to collect the `tech-support` through management IP, defaults to all hosts.
- **exclude-mgmt-node**: It is an option not to collect `tech-support` from the management node.

Tech-Support Configuration File

Cisco VIM `tech-support` is a utility tool is designed to collect the VIM pod logs which help users to debug the issues offline. The administrator uses the `tech-support` configuration files to provide the list of commands or configuration files. The `tech support` tool of the Cisco VIM gathers list of commands or configuration files for the offline diagnostic or debugging purposes.

By default the `tech-support` configuration file is located at the `/root/openstack-configs/tech-support/tech_support_cfg.yaml` file. Alternatively, you can use a different one by specifying the `-config-file` option. The syntax of this configuration file must be as follows:

The `tech-support` configuration file section is divided into eight sections which corresponds to each of the installer stages:

- `INPUT_VALIDATION`
- `BUILDNODE_ORCHESTRATION`
- `VALIDATION`
- `BAREMETAL_INSTALL`
- `COMMON_SETUP`
- `CEPH`
- `ORCHESTRATION`
- `VMTP`

Inside each of these eight sections, there are tags divided on hierarchical levels. At the first level, the tag indicates the host(s) or path on which the command(s) run and from where the file(s) can be collected. The possible tags are as follows:

- - HOSTS_MANAGEMENT: Run in the Management node only
- - HOSTS_CONTROL: Run in all the Control nodes
- - HOSTS_COMPUTE: Run in all the Compute nodes
- - HOSTS_STORAGE: Run in all the Storage nodes
- - HOSTS_COMMON: Run in all the Compute and Control nodes
- - HOSTS_ALL: Run in all the Compute, Control and Storage nodes



Note In any of these eight sections, if HOSTS tag is not specified then no information is collected for that stage.

For each of the hosts mentioned above there is a second level tag which specifies where to run the command. The possible values of those tags are as follows:

- - SERVER_FILES: Path(s) to the file(s) that tech-support has to collect.
- - SERVER_COMMANDS: Command(s) or script name(s) which has to be executed directly on the server. The command(s) has to be included before in the \$PATH. For the scripts, refer to the Custom Scripts paragraph below.
- - CONTAINERS: Indicates the tech-support tool that the command(s) has to be executed and the files to be gathered from inside a container. See the following steps for more specific information of what can be added in this section.

In the CONTAINERS section, indicate the path in which container the commands are executed or gathered from. This is done with a <container_name> tag The following are the shown to get the string for the <container_name> tag):

- all_containers: Execute inside all containers (regardless of the state).
- <container_name>: Container Name must be the name of a container and it indicates in which container to run the command or gather the information. It runs commands inside the container only if the mentioned container is up (as we cannot run commands on dead containers). Examples of how to get the container name:
 - Execute **docker ps** and get the name (without any numbers) of the last column of output **docker ps -a**.

For example:

```
CONTAINER ID  IMAGE  COMMAND  <snip>  NAMES
81bc4e54cbfb  <registry>/vmtp:4263  /bin/bash"  vmtp_4263
```

The tech-support runs the linux commands on the server (from packages that is included in RHEL7.3). Add the name of the commands under the SERVER_COMMANDS section of the configuration file to run the commands.

However, if the administrator wants to add a custom bash or python script to be executed in some set of servers in the cloud. In such case you need to add the script into the custom-scripts directory on the current directory path (/root/openstack-configs/tech-support/) and add the script name into the corresponding SERVER_COMMANDS section.

The tech-support tool will scp the script(s) included in the custom-scripts directory into the appropriate cloud nodes where it will be executed (as# indicated in this config file) and capture the output (stdout and stderr) and add it to the collection of files collected by the tech-support tool. It is assumed that the scripts are self-standing and independent and needs no external input.

Following is an example of a custom tech-support configuration file. This is just an example of what information the tech-support tool will gather if given the following configuration file:

```
COMMON_SETUP:
HOSTS_ALL: # All compute, control and storage hosts
SERVER_FILES:
- /usr/lib/docker-storage-setup
SERVER_COMMANDS:
- docker info
- my_script.sh
CONTAINERS:
all_containers: #execute in all containers (even if they are in down state)
CONTAINER_COMMANDS:
- docker inspect
- docker logs
logstash:
CONTAINER_FILES:
- /var/log/
CONTAINER_COMMANDS:
- ls -l
```

Given this example of configuration, and assuming that the installer ended in at least the COMMON_SETUP state, the tech-support tool will run under all OpenStack nodes (Compute, Control and Storage) and it will:

- Gather (if exists) the contents of /usr/lib/docker-storage-setup file.
- Run **docker info** command and collect the output.
- Run **my_script.sh** and collect the output. The **my_script.sh** is an example of a bash script which the user previously added to the /root/openstack-configs/tech-support/custom-scripts directory.
- Collect the output of docker inspect and docker logs for all containers.
- Collect the files in /var/log inside the logstash container (if there is container with that name). This is equivalent to running the following command (where /tmp indicates a temporary location where the tech-support tool gathers all the information): **docker cp logstash_{tag}:/var/log/ /tmp**.
- Collect the output of the command **docker exec logstash_{{tag}}: ls -l**.

Tech-Support When Servers Are Offline

It is difficult to collect the information from the servers if one or more cloud nodes are not reachable. In this case, you can connect through the KVM console into those servers and run the local tech-support tool.

Step 1 To run the local tech-support tool run the following command:

```
/root/tech_support_offline
```

Step 2 Cisco VIM tech_support_offline collects the Logs and other troubleshooting output from the server and place it in the location of the other server:

```
/root/tech_support
```

Note After the server is reachable, you can use the Cisco VIM tech-support tool which collects all the files under the /root/tech-support/ directory which can be used to debug any issue which are offline.

Disk-Maintenance Tool to Manage Physical Drives

In VIM you can use the disk-maintenance tool to check the status of all physical drives that are present in running and operational nodes in the following roles -

- Management
- Control (all or specific nodes)
- Compute (all or specific nodes) (Expect for third party)

This provides the information about the present status of the physical drives - if they are in Online, Offline, Rebuilding, Unconfigured Good or JBOD states if all disks are ok. If not, the disks that have gone bad are displayed with the slot number and server information, that has to be replaced. When multiple disks have to be replaced, we recommend you to execute remove or add of the node.

- Physically remove and insert a new disk before attempting to replace.
- For smooth operation, wipe out disk before attempting replace operations.
- Call Cisco TAC if you face any issue. Do not reattempt.



Note Make sure that each node is running with hardware RAID, the steps for which can be found in the section titled Recovering Compute Node Scenario. Refer to step 15 of the section "Upgrading Cisco VIM Software Using a USB" on how to move the pod from hardware RAID to software RAID.

To check the status of the Diskmgmt log in to the management node and run the ciscovim command with the diskmgmt option. The design of the diskmgmt user interface follows a test job create, list, show, and delete workflow.

Diskmgmt user workflow:

A database of disk operation results is maintained so that you can keep the results of multiple disk check or replace and view them at any time.

Step 1 Run the Help command to see all available command line options:

```
# ciscovim help diskmgmt
usage: ciscovim diskmgmt [--server <node1,node2,...>] [--id <id>]
```

```
[--locator {on,off}] [--json-display] [-y]
create|delete|list|show check-disks|replace-disks
all|management|control|compute
```

HDD maintenance helper

Positional arguments:

```
create|delete|list|show      The control command to perform
check-disks|replace-disks   The identity of the task/action
all|management|control|compute The role of the target host(s)
```

Optional arguments:

```
--server <node1,node2,...> List of specific control/compute host names
                              within the target role.
--id <id>                    ID used to identify specific item to
                              show/delete.
--locator {on,off}          Turn on/off locator LED for server with bad
                              disks and for the physical drives.
--json-display              Shows output in JSON format.
-y, --yes                   Yes option to perform the action
```

Step 2 Check Disk operation creates check-disks operation for all control nodes in the POD. The system responds with a message indicating the Time, ID and when it was Created. Run the following check-disk operation command:

```
# ciscovim diskmgmt create check-disks control
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| action     | check-disks                             |
| command    | create                                   |
| created_at | 2018-03-07T21:12:20.684648+00:00        |
| id         | 0c6d27c8-bdac-493b-817e-1ea8640dae57    |
| locator    | False                                    |
| result     |                                           |
| role       | control                                  |
| servers    | None                                      |
| status     | not_run                                  |
| updated_at | None                                      |
+-----+-----+
```

Step 3 The cisco vim diskmgmt list command is used to monitor a currently running task, and the completed tasks. The list command can filter based on the role. Using ‘all’ command lists all tests that are in the database.

```
# ciscovim diskmgmt list check-disks control
+-----+-----+-----+-----+-----+
| ID                | Action      | Role    | Status  | Created                |
+-----+-----+-----+-----+-----+
| 861d4d73-ffee-40bf-9348-13afc697ee3d | check-disks | control | Complete | 2018-03-05 14:44:47+00:00 |
| 0c6d27c8-bdac-493b-817e-1ea8640dae57 | check-disks | control | Running  | 2018-03-07 21:12:20+00:00 |
+-----+-----+-----+-----+-----+
[root@F24-Michigan ~]# ciscovim diskmgmt list check-disks compute
+-----+-----+-----+-----+-----+
| ID                | Action      | Role    | Status  | Created                |
+-----+-----+-----+-----+-----+
| 0be7a55a-37fe-43a1-a975-cbf93ac78893 | check-disks | compute | Complete | 2018-03-05 14:45:45+00:00 |
+-----+-----+-----+-----+-----+
[root@F24-Michigan ~]# ciscovim diskmgmt list check-disks all
+-----+-----+-----+-----+-----+
| ID                | Action      | Role    | Status  | Created                |
```

```

+-----+-----+-----+-----+
| cdfd18c1-6346-47a2-b0f5-661305b5d160 | check-disks | all | Complete | 2018-03-05 14:43:50+00:00 |
| 861d4d73-ffee-40bf-9348-13afc697ee3d | check-disks | control | Complete | 2018-03-05 14:44:47+00:00 |
| 0be7a55a-37fe-43a1-a975-cbf93ac78893 | check-disks | compute | Complete | 2018-03-05 14:45:45+00:00 |
| 0c6d27c8-bdac-493b-817e-1ea8640dae57 | check-disks | control | Complete | 2018-03-07 21:12:20+00:00 |
+-----+-----+-----+-----+

```

Step 4 Run the following command to show the detailed results of a diskmgmt check-disks operation:

```
# ciscovim diskmgmt show check-disks control --id 0c6d27c8-bdac-493b-817e-1ea8640dae57
```

```

+-----+-----+-----+-----+-----+
| Message | Host | Role | Server | State |
+-----+-----+-----+-----+-----+
| Raid Health Status | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | Optimal |
| | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | Optimal |
| | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | Optimal |
| | | | | |
| VD Health Status | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | Optimal |
| | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | Optimal |
| | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | Optimal |
| | | | | |
| RAID Level and Type | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | Type - HW; Level - RAID1 |
| | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | Type - HW; Level - RAID1 |
| | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | Type - HW; Level - RAID1 |
| | | | | |
| Number of Physical Disks | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | 8 |
| | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | 8 |
| | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | 8 |
| | | | | |
| Number of Virtual Disks | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | 1 |
| | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | 1 |
| | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | 1 |
| | | | | |
| Boot Drive Disk Media-Type | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | HDD |
| | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | HDD |

```



```
[--json-display] [-y]
create|delete|list|show check-osds|replace-osd
```

OSD maintenance helper

Positional arguments:

```
create|delete|list|show      The control command to perform
check-osds|replace-osd     The identity of the task/action
```

Optional arguments:

```
--server <node1,node2,...> List of specific block_storage hostnames
--detail                    Display full OSD details
--id <id>                   ID used to identify specific item to
                             show/delete.
--osd <osd_name>           Name of down OSD to replace. Eg. 'osd.xx'
--locator {on,off}         Turn on|off locator LED for server with bad OSDs
                             and for the physical drives.
--json-display              Show output will be in JSON format.
-y, --yes                   Yes option to perform the action
```

```
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Step 2 To check the osds run the following command:

```
# ciscovim osdmgmt create check-osds
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Field      | Value                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| action     | check-osds                              |
| command    | create                                   |
| created_at | 2018-03-08T21:11:13.611786+00:00        |
| id         | 5fd4f9b5-786a-4a21-a70f-bffac70a3f3f   |
| locator    | False                                    |
| osd        | None                                     |
| result     |                                           |
| servers    | None                                     |
| status     | not_run                                  |
| updated_at | None                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Step 3 Monitor the osdmgmt check operations using the list command. Cisco Vim Osd mgmt list commands are used to monitor the currently running test. It also helps you to view the tests that are run/ completed.

```
# ciscovim osdmgmt list check-osds
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID                                               | Action   | Status   | Created                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 5fd4f9b5-786a-4a21-a70f-bffac70a3f3f          | check-osds | Complete | 2018-03-08 21:11:13+00:00                 |
| 4efd0be8-a76c-4bc3-89ce-142de458d844          | check-osds | Complete | 2018-03-08 21:31:01+00:00                 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Step 4 To show the detailed results of a osdmgmt check-osds operation run the following command:

```
# ciscovim osdmgmt show check-osds --id 5fd4f9b5-786a-4a21-a70f-bffac70a3f3f
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Message                | Host                | Role                | Server  | State |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Overall OSD Status    | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | Optimal |
|                       | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | Optimal |
|                       | f24-michigan-micro-2 | block_storage control compute | 7.7.7.6 | Optimal |
|                       |                       |                       |         |       |
| Number of OSDs        | f24-michigan-micro-3 | block_storage control compute | 7.7.7.5 | 5       |
|                       | f24-michigan-micro-1 | block_storage control compute | 7.7.7.7 | 5       |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```


Utility to Resolve Cisco VIM Hardware Validation Failures

The Cisco VIM Hardware Validation utility tool is used to perform hardware validation during the installation of UCS C-series servers. It captures the user and environmental hardware validation errors that occur during the installation process. The tool enables you to fix these errors that are based on the inputs you provide at the Command Line Interface (CLI). It validates the updated configurations to verify if the changes are applied properly. After the error is resolved, you can resume the installation from the point of failure.

The `ciscovim hardware-mgmt` user interface allows you to test the job `validate orresolve-failures(create)`, `list`, `show`, and `delete` workflow

Hardware-mgmt user workflow:

1. Use “`ciscovim hardware-mgmt validate ...`” command to initiate a validation.
2. Use “`ciscovim hardware-mgmt list ...`” command to view summary/status of current test jobs.
3. Use “`ciscovim hardware-mgmt show ... --id <ID>`” command to view detail test results
4. Use “`ciscovim hardware-mgmt delete ... --id <ID>`” to delete test results.

A database of results is maintained so that the user can keep the results of multiple hardware-mgmt operations and view them at any time.



Note You cannot use the utility for the following tasks:

- Configuring BIOS settings for the B-series pods.
- Upgrading or changing the firmware version.
- Resolving hardware failures other than `lom`, `hba`, `flexflash`, `pcie_slot`, `power`, and `vnuc_pxe_boot`.

Command Usage

To capture the list of failures that can be resolved by using the utility, go to the `install` directory and execute the `help` command:

```
# cd <installer-id>/clouddeploy
```

```
# python hw_validations.py -help .
```

The following shows the output of the `help` command.

```
usage: hw_validations.py [-h] [--resolve-failures RESOLVE_FAILURES]
  [--validate VALIDATE_OF] [-y] [--host HOSTS]
  [--file SETUP_FILE_LOCATION]
UCS Hardware Validations
optional arguments:
-h, --help show this help message and exit
--resolve-failures RESOLVE_FAILURES, -rf RESOLVE_FAILURES
    all - Fix all the failures.
    lom - Fix LOM port(s) status failures.
    hba - Fix HBA port status failures.
    flexflash - Fix Flexflash failures.
```

```

pcie_slot - Fix PCIe slot status failures.
power - Fix Power failures.
vnic_pxe_boot - Fix Vnic PXE_Boot statusfailures
-y, -yes
--host HOSTS Comma separated list of hostnames
--file SETUP_FILE_LOCATION, -f SETUP_FILE_LOCATION
    Provide a valid 'setup_data.yaml' file
    
```

Command Syntax

hw_validations.py [-h] [--resolve-failures RESOLVE_FAILURES] [--validate VALIDATE_OF] [-y] [--host HOSTS] [--file SETUP_FILE_LOCATION]

The following table provides the description of the parameters of the command.

Optional	Description
[-h], --help	Provides detailed information about the command.
[--resolve-failures RESOLVE_FAILURES], -rf RESOLVE_FAILURES	Enables you to specify the failure that you want to resolve. The optional arguments are as follows:
[-y]	Yes
[--host HOSTS]	Enables you to specify the hostname of the server for which you want to resolve failures. You cannot specify the IP address or CIMC IP address of servers as arguments. You can specify a list of hostnames as comma-separated arguments. If the -host option is not specified, the failures of all the servers that are specified in the setup_data.yaml file are resolved.
[--file SETUP_FILE_LOCATION] [-f SETUP_FILE_LOCATION]	Enables you to specify the name of a setup_data.yaml file.

Examples of Command Usage

The following table provides the commonly used commands along with their examples.

Purpose	Syntax	Example
To resolve all failures.	python hw_validations.py --resolve-failures all -y	python hw_validations.py --resolve-failures all -y
To simultaneously resolve one or more failures.	python hw_validations.py --resolve-failures <failure-1>,<failure-2> -y	To resolve the lom and hba status failures: python hw_validations.py --resolve-failures lom,hba -y

Purpose	Syntax	Example
To resolve the errors by using the setup_data.yaml file.	python hw_validations.py --resolve-failures <failure-1>,<failure-2> -y --file <location-of-yaml file>	To resolve the LOM status failures by using ~/save/setup_data.yaml file: python hw_validations.py --resolve-failures lom,hba -y --file ~/save/setup_data.yaml
To resolve failures on a particular server as specified in the setup_data.yaml file by using the -- host option.	python hw_validations.py --resolve-failures <failure-1> -y --host <name-of-host-server-1>,<name-of-host-server-2>	To resolve the PCIe slot failures on hiccup-controller-1 server as specified in the setup_data.yaml: python hw_validations.py --resolve-failures pcie_slot -y --host hiccup-controller-1

Cisco VIM Client Debug Option

The --debug option enables you to get verbose logging on the ciscovim client console. You can use verbose logging to troubleshoot issues with the ciscovim client.

The debug option has the following parts:

- Curl Command: Curl command can be used for debugging. It can be executed standalone. Curl Command also displays the REST API Endpoint and the Request Payload.
- Response of REST API

Examples of Using debug Option to get list of passwords

```
# ciscovim --debug list-password-keys
2018-05-28 22:13:21,945 DEBUG [ciscovimclient.common.httpClient][MainThread] curl -i -X GET
-H 'Content-Type: application/json' -H 'Authorization: ****' -H 'Accept: application/json'
-H 'User-Agent: python-ciscovimclient' --cacert /var/www/mercury/mercury-ca.crt
https://172.31.231.17:8445/secrets
2018-05-28 22:13:21,972 DEBUG [ciscovimclient.common.httpClient][MainThread]
HTTP/1.1 200 OK
content-length: 1284
x-xss-protection: 1
x-content-type-options: nosniff
strict-transport-security: max-age=31536000
server: WSGIServer/0.1 Python/2.7.5
cache-control: no-cache, no-store, must-revalidate, max-age=0
date: Tue, 29 May 2018 05:13:21 GMT
x-frame-options: SAMEORIGIN
content-type: application/json; charset=UTF-8

{u'HEAT_KEYSTONE_PASSWORD': '****', u'CINDER_KEYSTONE_PASSWORD': '****',
u'METADATA_PROXY_SHARED_SECRET': '****', u'WSREP_PASSWORD': '****', u'ETCD_ROOT_PASSWORD':
'****', u'HEAT_DB_PASSWORD': '****', u'CINDER_DB_PASSWORD': '****', u'KEYSTONE_DB_PASSWORD':
'****', u'NOVA_DB_PASSWORD': '****', u'GLANCE_KEYSTONE_PASSWORD': '****',
u'CLOUDPULSE_KEYSTONE_PASSWORD': '****', u'VPP_ETCD_PASSWORD': '****', u'COBBLER_PASSWORD':
'****', u'DB_ROOT_PASSWORD': '****', u'NEUTRON_KEYSTONE_PASSWORD': '****',
u'HEAT_STACK_DOMAIN_ADMIN_PASSWORD': '****', u'KIBANA_PASSWORD': '****',
u'IRONIC_KEYSTONE_PASSWORD': '****', u'ADMIN_USER_PASSWORD': '****', u'HAPROXY_PASSWORD':
'****', u'NEUTRON_DB_PASSWORD': '****', u'IRONIC_DB_PASSWORD': '****', u'GLANCE_DB_PASSWORD':
```

```
'****', u'RABBITMQ_ERLANG_COOKIE': '****', u'NOVA_KEYSTONE_PASSWORD': '****',
u'CPULSE_DB_PASSWORD': '****', u'HORIZON_SECRET_KEY': '****', u'RABBITMQ_PASSWORD': '****'}
```

```
+-----+
| Password Keys |
+-----+
| ADMIN_USER_PASSWORD |
| CINDER_DB_PASSWORD |
| CINDER_KEYSTONE_PASSWORD |
| CLOUDPULSE_KEYSTONE_PASSWORD |
| COBBLER_PASSWORD |
| CPULSE_DB_PASSWORD |
| DB_ROOT_PASSWORD |
| ETCD_ROOT_PASSWORD |
| GLANCE_DB_PASSWORD |
| GLANCE_KEYSTONE_PASSWORD |
| HAPROXY_PASSWORD |
| HEAT_DB_PASSWORD |
| HEAT_KEYSTONE_PASSWORD |
| HEAT_STACK_DOMAIN_ADMIN_PASSWORD |
| HORIZON_SECRET_KEY |
| IRONIC_DB_PASSWORD |
| IRONIC_KEYSTONE_PASSWORD |
| KEYSTONE_DB_PASSWORD |
| KIBANA_PASSWORD |
| METADATA_PROXY_SHARED_SECRET |
| NEUTRON_DB_PASSWORD |
| NEUTRON_KEYSTONE_PASSWORD |
| NOVA_DB_PASSWORD |
| NOVA_KEYSTONE_PASSWORD |
| RABBITMQ_ERLANG_COOKIE |
| RABBITMQ_PASSWORD |
| VPP_ETCD_PASSWORD |
| WSREP_PASSWORD |
+-----+
```

Examples of Using debug option to get list of nodes

```
# ciscovim --debug list-nodes
2018-05-28 22:13:31,572 DEBUG [ciscovimclient.common.httpclient][MainThread] curl -i -X GET
-H 'Content-Type: application/json' -H 'Authorization: ****' -H 'Accept: application/json'
-H 'User-Agent: python-ciscovimclient' --cacert /var/www/mercury/mercury-ca.crt
https://172.31.231.17:8445/nodes
2018-05-28 22:13:31,599 DEBUG [ciscovimclient.common.httpclient][MainThread]
HTTP/1.1 200 OK
content-length: 2339
x-xss-protection: 1
x-content-type-options: nosniff
strict-transport-security: max-age=31536000
server: WSGIServer/0.1 Python/2.7.5
cache-control: no-cache, no-store, must-revalidate, max-age=0
date: Tue, 29 May 2018 05:13:31 GMT
x-frame-options: SAMEORIGIN
content-type: application/json; charset=UTF-8

{'nodes': {'status': 'Active', 'uuid': 'u'6b1ea6ee-b15b-41ca-9d79-3bb9ec0002bc',
'setupdata': 'u'fe78b5f9-5a46-447c-9317-2bf7362c1e81', 'node_data': {'rack_info':
{'rack_id': 'u'RackD'}, 'cimc_info': {'cimc_ip': 'u'172.29.172.81'}, 'management_ip':
'u'21.0.0.10'}, 'updated_at': 'u'2018-05-25T11:14:46+00:00', 'reboot_required': 'u'No',
'mtype': 'u'control', 'install': 'u'372aa3c1-1ab0-4dd0-a8a8-1853a085133c', 'power_status':
'u'PowerOnSuccess', 'install_logs':
u'https://172.31.231.17:8008/edd3975c-8b7c-4d3c-93de-a033ae10a6b6', 'created_at':
'u'2018-05-21T13:25:50+00:00', 'name': 'u'gg34-2'}}

+-----+-----+-----+-----+
```

Node Name	Status	Type	Management IP
gg34-1	Active	control	21.0.0.12
gg34-2	Active	control	21.0.0.10
gg34-3	Active	control	21.0.0.11
gg34-4	Active	compute	21.0.0.13

Example of Getting Response from REST API using Curl Commands

Get the REST API Password.

```
# cat /opt/cisco/ui_config.json
{
"Kibana-Url": "http://172.31.231.17:5601",
"RestAPI-Url": "https://172.31.231.17:8445",
"RestAPI-Username": "admin",
"RestAPI-Password": "*****",
"RestDB-Password": "*****",
"BuildNodeIP": "172.31.231.17"
}
```

Form the Curl Command.

```
curl -k -u <RestAPI-Username>:<RestAPI-Password> <RestAPI-Url>/<Endpoint>
```

E.g. To get Nodes Info of Cloud

```
curl -k -u admin:**** http://172.31.231.17:5601/v1/nodes
```

Examples of Response of REST APIs

API "/"

```
# curl -k -u admin:**** https://172.31.231.17:8445/
```

```
{
"default_version": {
"id": "v1",
"links": [
{"href": "http://127.0.0.1:8083/v1/",
"rel": "self"}
]
},
"versions": [
{
"id": "v1",
"links": [
{"href": "http://127.0.0.1:8083/v1/",
"rel": "self"}
]
}
]
,
"name": "Virtualized Infrastructure Manager Rest API",
"description": "Virtualized Infrastructure Manager Rest API is used to invoke installer from API."
}
```

API "/v1/setupdata/"

```
# curl -k -u admin:**** https://172.31.231.17:8445/v1/setupdata/
```

```
{
"setupdatas": [
. . .
]
}
```

API "/v1/nodes"

```
# curl -k -u admin:**** https://172.31.231.17:8445/v1/nodes
```

```
{
"nodes": [
{
"status": "Active",
"uuid": "0adabc97-f284-425b-ac63-2d336819fbaf",
"setupdata": "fe78b5f9-5a46-447c-9317-2bf7362c1e81",
"node_data": {
"rack_info": {
"rack_id": "RackC"},
"cimc_info": {
"cimc_ip": "172.29.172.75"},
"management_ip": "21.0.0.13"}
},
"updated_at": "2018-05-21T15:11:05+00:00",
"reboot_required": "No",
"mtype": "compute",
"install": "372aa3c1-1ab0-4dd0-a8a8-1853a085133c",
"power_status": "PowerOnSuccess",
"install_logs": "https://172.31.231.17:8008/edd3975c-8b7c-4d3c-93de-a033ae10a6b6",
"created_at": "2018-05-21T13:25:50+00:00",
"name": "gg34-4",
. . .
}
]
}
```

API "/v1/secrets"

```
# curl -k -u admin:**** https://172.31.231.17:8445/v1/secrets
```

```
{
"HEAT_KEYSTONE_PASSWORD": "5oNff4jWsvAwnWk1",
"CINDER_KEYSTONE_PASSWORD": "Hq4i6S5CnfQe7Z2W",
. . .
}
```

```

"RABBITMQ_ERLANG_COOKIE": "XRMHBQHTLVJSVWDFKJUX", "METADATA_PROXY_SHARED_SECRET":
"XNzrhosqW4rwiz7c", "WSREP_PASSWORD": "z1oQqhKd1fXDxJTV", "ETCD_ROOT_PASSWORD":
"LMLC8gvilIA3KiIc", "HEAT_DB_PASSWORD": "J8zt8ldMvdtJxAtG", "CINDER_DB_PASSWORD":
"BVX3y2280DSx2JkY", "KEYSTONE_DB_PASSWORD": "55fVNzxr1VxCNodh", "NOVA_DB_PASSWORD":
"RklMK1OIJgsjGZal", "IRONIC_KEYSTONE_PASSWORD": "9tYZgIw6SZERZ1dZ", "ADMIN_USER_PASSWORD":
"DjDQrk4QT7pgHy94", "GLANCE_KEYSTONE_PASSWORD": "w4REb8uhrHquCfRm", "HAPROXY_PASSWORD":
"oB0v7VJoo2IfB80W", "CLOUDPULSE_KEYSTONE_PASSWORD": "q6QVvxBQhrqv6ZhX", "NEUTRON_DB_PASSWORD":
"FZVMWgApcZR4us5q", "IRONIC_DB_PASSWORD": "dq3Udmu95DWyX1jy", "GLANCE_DB_PASSWORD":
"O7vQ2emuPDrvd4x", "KIBANA_PASSWORD": "azHHhP4ewxpZVwcg", "VPP_ETCD_PASSWORD":
"NLYIAvECMW2qI7Bp", "NOVA_KEYSTONE_PASSWORD": "JUfMNGz0BZG7JwXV", "NEUTRON_KEYSTONE_PASSWORD":
"QQ01o8Q87BjFoAYQ", "CPULSE_DB_PASSWORD": "DaFthNtpX2RvwTSS", "COBBLER_PASSWORD":
"XoIJ9mbWcmVyzvvN", "HORIZON_SECRET_KEY":
"NHkA0qwHIWUSwhPZowJ8Ge3RyRd6oM8XjOT8PHnZdckxgm3kbb1Msltsw0TAQJnx", "DB_ROOT_PASSWORD":
"seqh5DRIKP6ZsKJ8", "HEAT_STACK_DOMAIN_ADMIN_PASSWORD": "Vu6LexEadAxscsvY",
"RABBITMQ_PASSWORD": "LBoYoxuvGsMsl1TX"}

API "/v1/nodes/mgmt._node"

# curl -k -u admin:**** https://172.31.231.17:8445/v1/nodes/mgmt_node

{"api_ip": "172.31.231.17", "mgmt_ip": "21.0.0.2"}

```