



Cisco Nexus Fabric OpenStack Enabler Install Guide Version 1.0

First Published: March 16, 2015

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

Text Part Number: OL-32000-03

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

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

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

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

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

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

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: <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. (1110R)

© 2015 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface

Preface v

Audience v

Document Conventions v

Related Documentation for Cisco Nexus Fabric vi

Obtaining Documentation and Submitting a Service Request vii

CHAPTER 1

Hardware and Software Requirements 1

Requirements for Hardware and Software 1

CHAPTER 2

Prerequisite 3

Prerequisite for OpenStack 3

CHAPTER 3

OpenStack for Nexus Fabric Overview 5

Overview of OpenStack for Cisco Nexus Fabric 5

CHAPTER 4

Before You Start 9

Before Starting 9

CHAPTER 5

Servers Topology 11

Required Resources 11

Topology 11

CHAPTER 6

Nexus Fabric Enabler Installation 13

Nexus Fabric Enabler Installation 13

CHAPTER 7

Post Installation 17

Post Installation 17

CHAPTER 8

Configuration in the Leaf and Cisco Prime DCNM 19

Leaf Interface Configuration 19

DHCP Configuration in Cisco Prime DCNM 19

CHAPTER 9

Create Project and Launch VM 23

Create Project and Launch VM 23

Known Caveats 23

Steps to Create a Project 23

DCI Support 24

Steps to Create a User for the Project 24

Steps to Create the Network 24

Steps to Create Security Group 29

Steps to Launch the VM 30

CHAPTER 10

Scalability 31

Scalability 31

CHAPTER 11

Known Limitations and Caveats 33

Known Limitations and Caveats 33

CHAPTER 12

Technical Support Model 35

Technical Support Model 35



Preface

The Preface contains the following sections:

- [Audience, page v](#)
- [Document Conventions, page v](#)
- [Related Documentation for Cisco Nexus Fabric, page vi](#)
- [Obtaining Documentation and Submitting a Service Request, page vii](#)

Audience

This publication is for experienced network administrators who configure and maintain Cisco Nexus Fabric.

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

Convention	Description
[x {y z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
<i>variable</i>	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<code>boldface screen font</code>	Information you must enter is in boldface screen font.
<i><code>italic screen font</code></i>	Arguments for which you supply values are in italic screen font.
<>	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation for Cisco Nexus Fabric

The Cisco Dynamic Fabric Automation documentation is at the following URL: <http://www.cisco.com/c/en/us/support/cloud-systems-management/dynamic-fabric-automation/tsd-products-support-series-home.html>.

The Cisco Nexus 6000 Series documentation is at the following URL: <http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/tsd-products-support-series-home.html>.

The Cisco Nexus 7000 Series documentation is at the following URL: <http://www.cisco.com/c/en/us/support/switches/nexus-7000-series-switches/tsd-products-support-series-home.html>.

The Cisco Nexus 5500 and 5600 Series documentation is at the following URL: <http://www.cisco.com/c/en/us/support/switches/nexus-5000-series-switches/tsd-products-support-series-home.html>.

The Cisco Nexus 1000V switch for VMware vSphere documentation is at the following URL: http://www.cisco.com/en/US/products/ps9902/tsd_products_support_series_home.html. The documentation therein includes the following guides for Cisco DFA. Additional information pertaining to troubleshooting can be located in the Cisco Nexus 1000V documentation for Cisco NX-OS Release 4.2(1)SV2(2.2).

- *Cisco Nexus 1000V DFA Configuration Guide, Release 4.2(1)SV2(2.2)*
- *Cisco Nexus 1000V VDP Configuration Guide, Release 4.2(1)SV2(2.2)*

The Cisco Prime Data Center Network Manager (DCNM) documentation is at the following URL: http://www.cisco.com/en/US/products/ps9369/tsd_products_support_series_home.html. The Cisco Prime DCNM documentation for Cisco DFA includes but is not limited to the following guides:

- *Cisco DCNM 7.0 OVA Installation Guide.*
- *Cisco DCNM 7.0 Fundamentals Guide*
- *Cisco DCNM DFA REST 7.0 API Guide*

The Cisco Prime Network Services Controller (NSC) documentation is at the following URL: http://www.cisco.com/en/US/products/ps13213/tsd_products_support_series_home.html.

The OpenStack for Cisco DFA install documentation includes the following guide and documents:

- *Open Source Used In OpenStack for Cisco DFA 1.0* at the following URL: http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/opensource/OpenStack_for_Cisco_DFA_1.0_Open_Source_Documentation.pdf
- *OpenStack for Cisco DFA Install Guide Using Cisco OpenStack Installer* at the following URL: <http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/install/guide/os-dfa-coi.pdf>
- *OpenStack for Cisco DFA Install Guide for Using Pre-built OpenStack for Cisco DFA Images* at the following URL: <http://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/dfa/openstack/install/guide/preblt-image.pdf>
- *Quick Guide to Clonezilla* at the following URL: <http://www.cisco.com/en/US/docs/switches/datacenter/dfa/openstack/install/guide/clonezilla-image-restore.pdf>

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation*, at: <http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.



Hardware and Software Requirements

- [Requirements for Hardware and Software, page 1](#)

Requirements for Hardware and Software

Cisco Nexus Fabric Enabler is a set of software that interacts with OpenStack through its open API's to allow users to connect Cisco Nexus 5600 Platform Switches and Cisco Nexus 6000 Series Switches as the network to the OpenStack compute nodes to form a cloud. You need to run Cisco NX-OS 7.1(0)N1(1) or later for the switch image.

The uplink port of server is directly connected to the fabric leaf switch and it generally requires that there is no middle device as LLDP needs to run between the server port and leaf switch port to signal the VM presence to the fabric for reachability.

OpenStack Juno release October 2014, needs to be installed on any server using an OpenStack installer of your choice. This guide does not cover about OpenStack installation.

- The qualification for this release is based on an OpenStack installation using DevStack.
- The qualification information of using other major third party OpenStack installer will be announced via release notes.

Basically there are three ways in which an Openstack installation can be done:

- Use DevStack for installation. This is primarily for development and testing purposes and not for production level installation.
- Use a third-party supported OpenStack installation package.
- Use the OpenStack general installation guidelines.

In addition to OpenStack installation, the other required installations are OVS and LLDPad - both are open source free software, similar to OpenStack. The Link Layer Discovery Protocol (LLDP) is an agent daemon to support VDP. Follow the respective README notes to build and install the latest versions on all servers that will be used as compute nodes.

Note that since OpenStack and its associated other pieces of software for this purpose is open source software, it generally requires you install the various software on your target servers as they come from different sources.

Cisco Nexus Fabric Enabler can be downloaded from the website https://github.com/CiscoSystems/fabric_enabler.



Prerequisite

- [Prerequisite for OpenStack, page 3](#)

Prerequisite for OpenStack

Following are the prerequisites for installing Nexus fabric OpenStack Enabler:

- Linux installation on your target servers
- Build and install LLDPad
- Build and install OVS

They need to be compiled and installed on all target compute nodes in the OpenStack setup as binaries could not be directly provided and used.

Linux installation on your target servers

We used Ubuntu 14.04.1 LTS with OpenStack Juno release in our development and QA environments via DevStack installer.

Build and install LLDPad

Use LLDPad open source community mailing list <http://open-lldp.org/> for your general queries.

- Clone the repository used for installation:
[git clone git://open-lldp.org/open-lldp](https://github.com/open-lldp/open-lldp)
- You do not need the Linux kernel installation, just the application



Note

All servers that will act as compute nodes (that is, where a VM can be launched by OpenStack) will need LLDPad installation.

Build and install OVS

You can skip this section if the installation of OVS in your setup is based on any release after April 24, 2014.

- Download OVS according to the instructions given in <http://www.openvswitch.org/download/>
- Use the most recent release or any stable release after April 24, 2014
- Follow the instructions in README carefully to build and install



Note

All servers that act as compute nodes will need OVS installation. Cisco QA uses openvswitch-2.3.1 for verification.



OpenStack for Nexus Fabric Overview

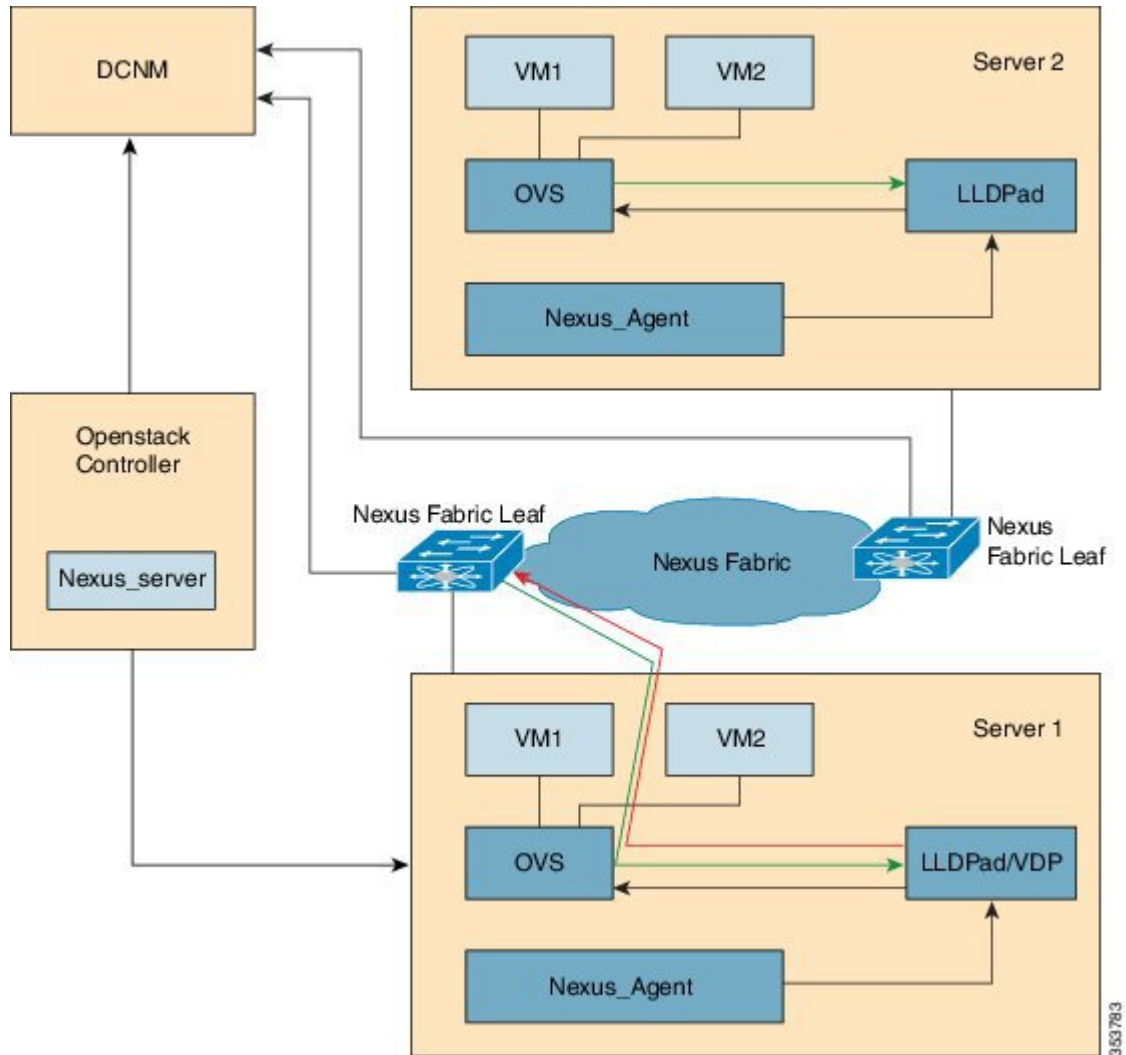
- [Overview of OpenStack for Cisco Nexus Fabric, page 5](#)

Overview of OpenStack for Cisco Nexus Fabric

OpenStack serves as one of orchestrators of the cloud enabled through Cisco Nexus Fabric. For this release, all the orchestration is done using OpenStack's dashboard Horizon graphic user interface or OpenStack generic open APIs.

The following diagram provides an overview of the system in the OpenStack with Cisco Nexus Fabric as its network connecting all the compute nodes:

Figure 1: System Flow



The control flow can be summarized as follows:

- 1 OpenStack cloud administrator creates a tenant/project. OpenStack Keystone component notifies Cisco Nexus Server running in the control node about the creation of the tenant. The Cisco Nexus Server sends the tenant information to Cisco Prime DCNM via Cisco Prime DCNM REST API. Tenant's user name and password are created by the administrator as well.
- 2 Tenant logs into OpenStack and create network. OpenStack Neutron components notifies Cisco Nexus Server of the Network Information (Subnet/mask, tenant name), and Cisco Nexus Server automatically sends the network creation request to Cisco Prime DCNM via Cisco Prime DCNM REST APIs.
- 3 Instance (VM) is launched, specifying the Network that the instance will be a part of.

- 4 Cisco Nexus Server sends the network information (Subnet/mask, tenant name) and VM information to Cisco Nexus agent running in the corresponding compute node.
- 5 VSI Discovery and configuration protocol (VDP), running in the context of LLDPad, gets notified by the Cisco Nexus agent about the VM and the segment ID associated with the VM.
- 6 VDP communicates with the Leaf passing the VM's information along with the Segment ID.
- 7 Leaf contacts Cisco Prime DCNM with the Segment ID for retrieving the Network attributes and auto-configure the compute node facing interface for this tenant VRF.
- 8 Leaf responds back with the VLAN to be used for tagging the VM's traffic. The VLAN to be used is the first free VLAN from the predefined VLAN pool (Cisco NX-OS command **system fabric dynamic-vlans**) on each individual leaf switch. The selected one is the leaf switch significance resource.
- 9 LLDPad module notifies the Cisco Nexus agent with the VLAN information provided by the leaf.

The Cisco Nexus agent configures OVS for tagging the packets from the VM destined to the network with the VLAN provided by the Leaf. VM's VNIC is operational only at this point.



Before You Start

- [Before Starting](#), page 9

Before Starting

- We assume you have some general understanding of Cisco Nexus Fabric system architecture. For more details, see <http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-fabric/white-paper-c11-731098.html>.
- We assume you know enough about OpenStack in general and have it installed and it is running in your setup (without L3 service, DHCP service or Neutron router configured).
- We assume you have each of OpenStack compute node connected to the Cisco Nexus switch leaf and OpenStack control node connected to Cisco Prime DCNM via an IP network. Of course, compute nodes and control node need to have their own connectivity per general OpenStack installation requirements. Cisco Prime DCNM 7.1 can be used with this OpenStack Enabler release.
- We assume you have installed LLDPad and OVS.
- Ensure that Cisco Prime DCNM is reachable via the Openstack Controller as well as from the Nexus Fabric cluster.
- The keystone and neutron services must run on the same server (the controller).
- All nodes must have same user account with the same credentials and also passwordless sudo has to be enabled on all nodes.

It can be checked in /etc/sudoers

```
%sudo ALL=(ALL:ALL) NOPASSWD: ALL
```

- The rabbitmq server must run on all the nodes. The rabbitmq is a general messaging scheme used by OpenStack software for passing information.



Servers Topology

- [Required Resources](#) , page 11
- [Topology](#), page 11

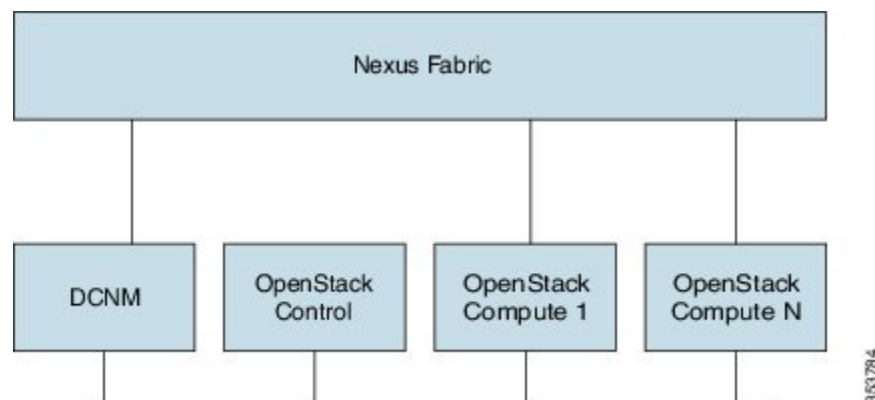
Required Resources

- DCNM needs to be running in the entire system. For more information, see [Cisco DCNM Installation Guide](#).
- One server is dedicated for the OpenStack Controller where the dashboard Horizon is running for doing orchestration. Optionally, control server node can also act as a compute node to host VM.
- Many compute nodes can be used in the setup, per general OpenStack scaling guidelines.

Topology

A sample topology is shown below. This is a critical step as it lays down your foundation so it is highly recommended that you follow the same wiring scheme.

Figure 2: Server Topology



The interface in the topology that is connected to the cluster should be operationally up (sudo ifconfig eth0 up).

The DCNM is also connected to the OpenStack control node via an IP network. You can optionally use the same server to run both control and compute node. In that case, you need to connect that server node to the Nexus Fabric leaf switch as well.



CHAPTER

6

Nexus Fabric Enabler Installation

- [Nexus Fabric Enabler Installation, page 13](#)

Nexus Fabric Enabler Installation

The sections below provide the information on installation of Nexus Fabric Enabler.

OpenStack Installation

For an OpenStack installation (see the three ways mentioned in the section above), specifically for supporting Nexus Fabric as the networking fabric for OpenStack, the [ml2] and [ovs] sections in *./etc/neutron/plugins/ml2/ml2_conf.ini* should be configured as follows:

```
[ml2]
type_drivers = local
mechanism_drivers = openvswitch

[ovs]
bridge_mappings = ethd:br-ethd

([ml2_type_flat], [ml2_type_vlan], [ml2_type_gre] and
[ml2_type_vxlan] sections should not be specified)
```

Ensure to perform the following at */etc/neutron/neutron.conf*:

```
core_plugin = neutron.plugins.ml2.plugin.ML2Plugin
```

Do not enable the following OpenStack services in an OpenStack installation:

- L3 agent
- DHCP service
- Do not set tunnel types

Install Nexus Fabric Enabler

Ensure your OpenStack installation is running correctly based on required setting discussed in the above sections.

On the control node in the OpenStack installation, get the source file

```
git clone -b enabler_1_0_0 https://github.com/CiscoSystems/fabric_enabler openstack_fabric_enabler
```

Go to the enabler directory created in the above step

```
cd ~/openstack_fabric_enabler
```

modify enabler_conf.ini

Following are the sections for modification in dcnm-mgmt interface (WebGui interface, Rest API and so on):

```
[dcnm]
dcnm_ip = 1.1.1.1
dcnm_user = username
dcnm_amqp_user = amqpdcnmusername
dcnm_password = password
gateway_mac = 20:20:00:00:00:AA
segmentation_id_min = segidmin
segmentation_id_max = segidmax
```

Ensure to replace the values `dcnm_ip`, `username`, `amqpdcnmusername`, `password`, `segidmin` and `segidmax` based on your setup. The `enhanced-fabric-mgmt (ssh, LDAP and so on)` command is introduced to solve the problem described in the `vswitch mobility` scenario. In the absence of this config, there will be greater traffic loss in the scenarios where a host move occurs for the hosts that were connected through `vswitch`.

**Note**

The user has to ensure that the segment ID should not be overlapping with other segment ID ranges (for example the Cisco Prime DCNM segment ID is using, default 30,000 to 49,999).

This is the section about Cisco Prime DCNM, which you have installed separately and set the right access credentials to be used by OpenStack Nexus Fabric Enabler. Ensure the `gateway_mac` matches your PoAP template setting in Cisco Prime DCNM for your leaf switch and you use the right range of segment IDs administrated by your fabric manager.

```
[dfa_mysql]
# MYSQL DB connection parameter
connection = mysql://username:password@localhost/cisco_dfa?charset=utf8
```

Ensure to replace the username and password based on your setting.

```
[dfa_rpc]
# Transport URL parameter for RPC
transport_url = 'rabbit://username:password@(ip)s:5672//'
```

Ensure to replace the username and password based on your setting.

```
[general]
compute_user = username
compute_passwd = password
```

Ensure to replace the username and password based on your setting.

Checking conf files

The `keystone_authtoken` in `neutron.conf` and `nova.conf` should be similar to the following setting:

For `neutron.conf`:

```
[keystone_authtoken]
signing_dir = /var/cache/neutron
auth_uri = http://<ip address of controller>:5000/v2.0
cafile = /opt/stack/data/ca-bundle.pem
identity_uri = http://<ip address of controller>:35357
auth_host = <ip address of controller>
auth_port = 35357
auth_protocol = http
```

```
admin_tenant_name = service
admin_user = neutron
admin_password = password
```

For nova.conf:

```
[keystone_authtoken]
signing_dir = /var/cache/nova
admin_password = password
admin_user = nova
admin_tenant_name = service
auth_uri = http://<ip address of controller>:5000/v2.0
cafile = /opt/stack/data/ca-bundle.pem
identity_uri = http://<ip address of controller>:35357
auth_protocol = http
auth_port = 35357
auth_host = <ip address of controller>
```



Note It requires all control/compute nodes that have same username/password and this is your Linux account on the servers to run as control/compute nodes.

Network Time Protocol (NTP)

The NTP installation is required by openstack (refer to openstack installation guide for NTP installation requirements).

In case of clock drift in openstack node, the following command can be executed to sync the node manually:

On ubuntu machine:

```
sudo service ntp stop
sudo ntpdate -s <your ntp server>
sudo service ntp start
```

On RedHat:

```
$ sudo vi /etc/sysconfig/ntpd
OPTIONS="-u ntp:ntp -p /var/run/ntpd.pid"
OPTION="-x -u ntp:ntp -p /var/run/ntpd.pid"
$ sudo service ntpd restart
```

Run python setup_enabler

Using the source file downloaded from [git clone https://github.com/CiscoSystems/fabric_enabler](https://github.com/CiscoSystems/fabric_enabler) openstack_fabric_enabler, the script will install all the necessary modules on both the control and compute nodes/servers and then bring them up.



Post Installation

- [Post Installation, page 17](#)

Post Installation

- Check if the Enabler server, agent and LLDPad are running

```
sudo status fabric_enabler_server (on node acting as control node only)
sudo status fabric_enabler_agent (on node acting as compute)
sudo status lldpad (on node acting as compute node)
```

- Check for existence of notification queues

```
sudo rabbitmqctl list_queues | grep cisco
cisco_dfa_keystone_notify.info 0
cisco_dfa_neutron_notify.info 0
```

- Check logs (the log file name and path is defined in the *enabler_conf.ini*) for any error

Add a Compute Node

When you add a new compute node to an existing OpenStack setup, the OpenStack installation for the new node needs to be done first. You can install Enabler to that node without affecting current OpenStack setup. Perform the following:

```
python setup_enabler --compute_name your-newly-added-compute-node-host-name
```




Configuration in the Leaf and Cisco Prime DCNM

- [Leaf Interface Configuration, page 19](#)
- [DHCP Configuration in Cisco Prime DCNM, page 19](#)

Leaf Interface Configuration

You must connect the leaf interface and accomplish the global configuration using the below step:

- Set the speed of the interface connected to compute node correctly, by entering the following. The **speed 1000** command is required when the server port of 1G is connected to Switch's 10G port.

```
fabric forwarding port-tracking
speed 1000
no shut
```

DHCP Configuration in Cisco Prime DCNM

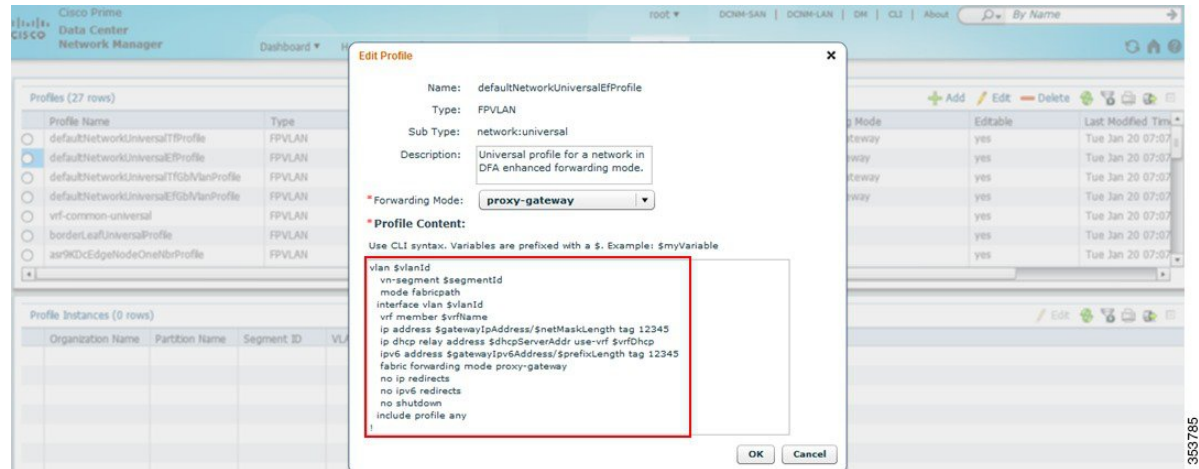
For DHCP configuration in Cisco Prime DCNM:

- 1 For auto configuration to work, the profile's DHCP address needs to be modified manually in Cisco Prime DCNM. This is because Openstack does not populate the DHCP address to Cisco Prime DCNM, since it has no way of knowing this information. For example, under a profile, you will see a config like, *ip dhcp relay address \$dhcpServerAddr use-vrf \$vrfDhcp*.

Modify this to the actual IP address of the DHCP server:

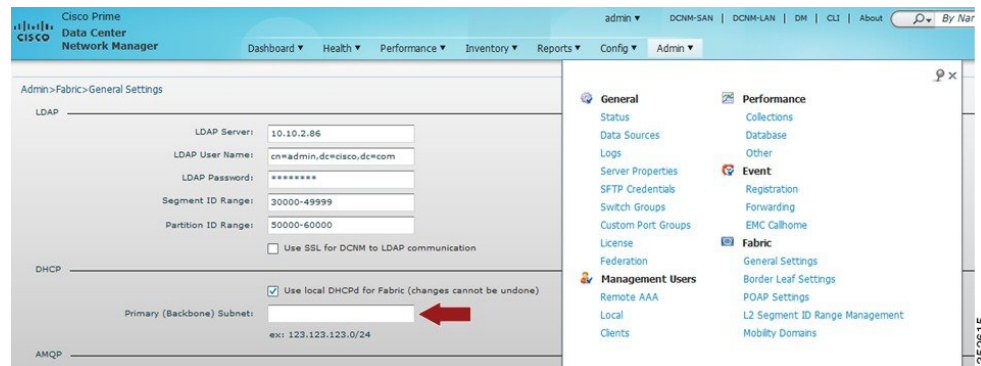
```
ip dhcp relay address 10.10.2.86 use-vrf management
```

Figure 3: Auto-Configuration Profile Tab



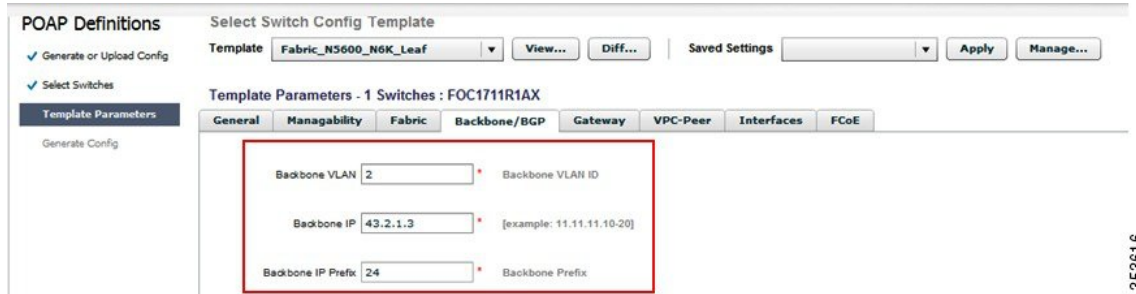
- 2 Fill the Nexus Fabric backbone VLAN subnet value in the Settings of the Cisco Nexus Fabric section under the Admin tab in the Cisco Prime DCM console.

Figure 4: Admin Tab



- 3 Ensure that Primary (Backbone) subnet with Use local DHCP for Cisco Nexus Fabric checked is consistent with the backbone VLAN subnet for the leaf and spine nodes in the Cisco Nexus Fabric. If the fabric is brought up from PoAP in Cisco Prime DCNM, the configuration screen will look like the following:

Figure 5: PoAP Definitions



353616



Create Project and Launch VM

- [Create Project and Launch VM, page 23](#)
- [Known Caveats, page 23](#)
- [Steps to Create a Project, page 23](#)
- [Steps to Create a User for the Project, page 24](#)
- [Steps to Create the Network, page 24](#)
- [Steps to Create Security Group, page 29](#)
- [Steps to Launch the VM, page 30](#)

Create Project and Launch VM

The information provided in this section is generic to OpenStack and it is provided here for your convenience with the exception of ConfigProfile, which is Cisco Fabric specific.

Known Caveats

Following are the known caveats:

- The Bulk Create and Delete functionality of VMs is not completely supported. (Refer the Known Limitations and Caveats section.)
- Do not give space for a network name or project name or a VM name.

Steps to Create a Project

Follow the steps to create a project:

- 1 Login to Horizon as admin, use the password you put into the OpenStack config file.
- 2 Click **projects** and then **create project**.

3 Click **Create Project**.

**Note**

The project name is used as the vrfName in the fabric (vrfName = "project_name:CTX") for fabric auto-configuration. The fabric limits the size of vrfName string to be 32 characters long currently. So make sure the project name length is less than 29 characters when creating the project. Do not use hyphens in the project name.

DCI Support

You can use OpenStack to configure DC Inter-connect. Note that we support only L3 DCI with Cisco Prime DCNM 7.1(1) release and Cisco NX-OS 7.1(0)N1(1) release or later.

As part of the project name string, type `xyx:dcid_id:129` to create it ("129" used here is just an example). Type `xyz` or `xyz:dcid_id:0` to remove the DCI support for this project.

The integer entered is the DCI ID. Cisco Prime DCNM uses it as an indication that user desires to auto-configure the border leaves with this VRF and extend to the DCI edge devices (s). A zero value is interpreted by Cisco Prime DCNM to remove the VRF configurations from border leaf and the configurations to extend the VRF from border leaf to DC Edge devices.

Steps to Create a User for the Project

Follow the steps to create a user for the project:

- 1 Click **Users** and then **Create User**.
- 2 Fill in all the fields, select the project you just created and select the role as 'admin'. The network information will not be populated to DCNM correctly, if you fail to specify the role as 'admin'.

Steps to Create the Network

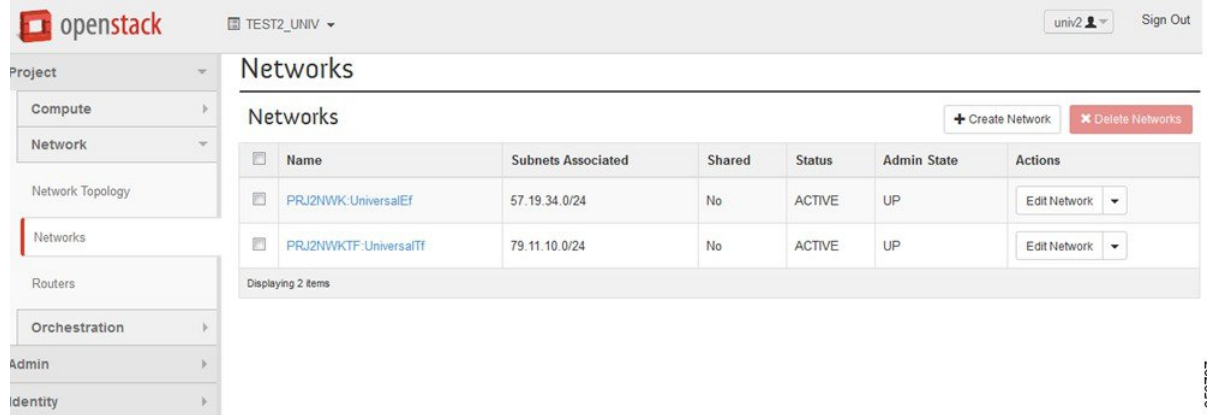
Follow the steps to create the network:

- 1 Login as the user created by the admin.
- 2 Click the **project** tab.
- 3 Click **Networks** and then click **Create Network**. Specify a Name for the network and go to the subnet tab. This is mandatory.
- 4 Specify a Network Address for the subnet.

Use non-default network profiles

By default for DCNM with version 7.1, defaultNetworkUniversalEfProfile will be the network profile automatically used by the system. Additionally defaultNetworkUniversalTfProfile can be specified too when creating network in OpenStack:

Figure 6: Default Network



Following are the supported network profiles with DCNM 7.1(1):

- defaultNetworkUniversalEfProfile
- defaultNetworkUniversalTfProfile
- defaultNetworkL2Profile

If it is an upgrade from version 7.0(1) or 7.0(2) to 7.1(1), the default profile will be defaultNetworkIpv4EfProfile and the supported profiles will be the sum of profiles for versions 7.0(1), 7.0(2) and 7.1(1) or later:

- defaultNetworkIpv4EfProfile
- defaultNetworkIpv4TfProfile
- defaultNetworkL2Profile
- defaultNetworkUniversalEfProfile
- defaultNetworkUniversalTfProfile
- defaultNetworkL2Profile

The syntax to specify the non-default profiles is the following when creating network in OpenStack by specifying the following name of network (“network_name” in the examples below) and a sub-string of the profile name:

- network_name:L2
- network_name: Ipv4Ef
- network_name: Ipv4Tf
- network_name: UniversalTf
- network_name: UniversalEf

Use defaultNetworkL2Profile

If this profile is chosen when a network is created in OpenStack, DCNM DHCP server will not assign IP address for the VM associated with this network. User is required to configure static IP address for the VM. Additionally, the following command needs to be run on the OpenStack control node:

```
$fabric_enabler_cli
Cisco Nexus Fabric Command Line Interface
(Nexus-Fabric) set_static_ip --mac fa:16:3e:72:ab:dc --ip 136.10.0.16
```

MAC address is of VNIC of the VM and IP address of the statically configured VM IP address. When a VM is removed from OpenStack, the above entry is automatically removed by the system.

Create networks from DCNM

In an OpenStack and Nexus Fabric powered cloud deployment, it is sometimes more desirable for some users to maintain management separation between servers and networks. This is possible by creating networks through DCNM, which is for network/fabric management primarily. First the project/tenant is created from OpenStack and then gets registered with DCNM. Then you can create network from DCNM.


Select Config and Auto-Configuration under Fabric as shown below:

Figure 7: Auto-Configuration



Then select the organization/partition in the following window:

Figure 8: Organization/Partition



Select dmz8_PR2 and click **Add** to add a network. Network created in DCNM can be deleted by selecting the network in the lower panel and click **Delete**. When performing **Add Network**, ensure Organization and Partition are the same as the project/tenant name in OpenStack and make sure the segment ID is in the same

range as set in the *enabler_conf.ini* file and it is not used by any existing networks. The *vrfName* with the correct format must be filled.

Figure 9: Add Network

Add Network

Organization: **dmz8_PR2**
 Partition: **CTX**
 VRF Name: dmz8_PR2:CTX
 Network Name:
 Network Role: **Host Network**
 Gateway IPv4 Address:
 Netmask Length:
 Gateway IPv6 Address:
 Prefix Length:

Network ID

Segment ID Only
 Segment ID:
 Mobility Domain and VLAN

DHCP Scope

Profile Name: **defaultNetworkUniversalTfProfile**

Profile Parameters
 Required 'vrfName' format is 'organizationName:partitionName'

vlanId:
 segmentId:
 vrfName: dmz8_PR2:CTX
 gatewayIpAddress:
 netMaskLength:
 dhcpServerAddr:
 vrfDhcp:

Service Configuration Parameters

VM Manager IP:
 Static IP Start:
 Static IP End:
 vSwitch Controller Network Id:
 Distributed Virtual Switch Id:
 Secondary Gateway IPv4 Address:

Also scroll down to add the DHCP range – this will make sure the VM will get DHCP IP address from this DCNM.


Figure 10: DHCP Scope

DHCP Scope

IP Range:

Click **OK** to save it. Log into OpenStack Horizon for that project/tenant and you can see network dmz6_net136 (2nd entry) is created with a suffix (DCNM) in the network name:

Figure 11: OpenStack



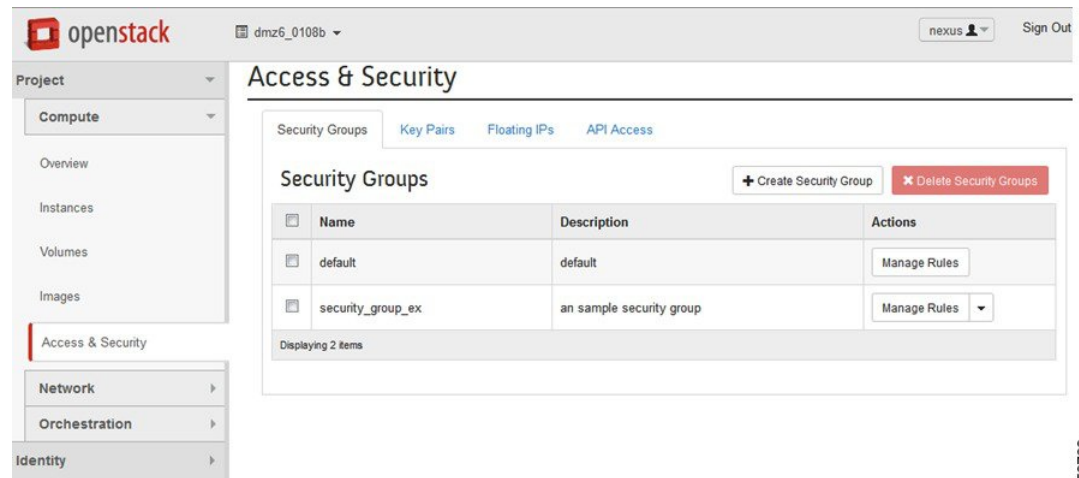
353792

This network is now available for you to associate when creating VMs. It is strongly recommended that the networks created by DCNM must be removed by DCNM after you remove all the VMs using this network.

Steps to Create Security Group

A security group with appropriate rules needs to be created and added before you can launch the VM. Create a security group and rules that allows DHCP (from DCNM) and your data traffic to go through. After logging into Horizon as a user, click **Project > Compute > Access Security**. Use the **Create Security Group** tab to create sample security group *security_group_ex* in the figure below.

Figure 12: Access and Security



353793

Click **Manage Rules** for the security group you just created and add new rules. For example, if the following rule is added for the security group created, it will allow all traffic:

Figure 13: Add Rule

Add Rule

Rule *
Other Protocol

Direction
Ingress

IP Protocol ⓘ
-1

Remote * ⓘ
CIDR

CIDR ⓘ
0.0.0.0/0

Description:
Rules define which traffic is allowed to instances assigned to the security group. A security group rule consists of three main parts:

Rule: You can specify the desired rule template or use custom rules, the options are Custom TCP Rule, Custom UDP Rule, or Custom ICMP Rule.

Open Port/Port Range: For TCP and UDP rules you may choose to open either a single port or a range of ports. Selecting the "Port Range" option will provide you with space to provide both the starting and ending ports for the range. For ICMP rules you instead specify an ICMP type and code in the spaces provided.

Remote: You must specify the source of the traffic to be allowed via this rule. You may do so either in the form of an IP address block (CIDR) or via a source group (Security Group). Selecting a security group as the source will allow any other instance in that security group access to any other instance via this rule.

353794

Steps to Launch the VM

Follow the steps to launch the VM:

- 1 Click **Instances** and then click **Launch Instance**.
- 2 Click **Image** drop-down menu and select the image.
There will be cirros image by default.
- 3 Specify a name for the Instance.
- 4 Select **Security** tab and choose the security group created (it is recommended to uncheck the default rule and select the one you specified).
- 5 Click **Networking** tab and select the network from the **Available network** list.
- 6 Click **Launch**.



Scalability

- [Scalability, page 31](#)

Scalability

OpenStack is supposed to support many compute nodes from the control node running the Horizon dashboard. This should not be related to the added Cisco Nexus Fabric functionality.

The recommended number of VMs per compute node for a user is ten.

Batching VM creation and deletion with OpenStack should be used with great caution. We support ten or fewer VMs per batch creation and deletion on two compute nodes. Batching capability beyond that is best-effort.



Known Limitations and Caveats

- [Known Limitations and Caveats, page 33](#)

Known Limitations and Caveats

The following are known limitations and caveats:

- Currently live-migration/vmotion is not supported.
- OpenStack does not support IPv6 yet.
- No service integration is supported.
- If a compute node with live VMs gets reloaded, OpenStack will lose these VMs.
- Currently a single port of a server running OpenStack is allowed to connect to a Fabric leaf port.
- The VMs' internal VNIC interface up/down is not detected by VDP. (Use VM deletion/creation as a potential workaround.)
- Servers with Converged Network Adapters (CNA) are not supported with OpenStack. The user needs to use a NIC.
- VM provisioned through Horizon will always get its IP address via DHCP server in DCNM.
- Cirros and mini Ubuntu-mini images are tested.
- Cisco UCS models with fabric interconnects are not supported currently with this release.
- Cisco Nexus 2000 series FEX support is currently not ready yet.



Technical Support Model

- [Technical Support Model](#), page 35

Technical Support Model

Following is the technical support model for using OpenStack:

- OpenStack is based on open source and is generally supported through its community using a best effort approach. If you use a third-party OpenStack installer and their support model, all generic OpenStack related support should come from the third-party support license.
- OpenStack Nexus Enabler part support will be from Cisco Nexus Solution Team.

