



Nexus Fabric Enabler Installation

- [Nexus Fabric Enabler Installation, page 1](#)

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.

