



OPEN
DAYLIGHT

OpenDaylight

and OpenStack

Beryllium (February 21, 2016)



OPEN
DAYLIGHT

OpenDaylight and OpenStack

OpenDaylight Community

Beryllium (2016-02-21)

Copyright © 2015 Linux Foundation All rights reserved.

This guide describes how to use OpenDaylight with OpenStack.

This program and the accompanying materials are made available under the terms of the Eclipse Public License v1.0 which accompanies this distribution, and is available at <http://www.eclipse.org/legal/epl-v10.html>

Table of Contents

Overview	5
1. Installing OpenStack	1
2. Installing OpenDaylight	2
OpenStack with OVSDB	2
OpenStack with GroupBasedPolicy	6
OpenStack with Virtual Tenant Network	12

List of Figures

2.1. Neutron Port	7
2.2. Neutron Network	8
2.3. Neutron Subnet	8
2.4. Neutron Security Group and Rules	9
2.5. Neutron Router	10
2.6. Neutron network mapping	11
2.7. Neutron subnet mapping	11
2.8. Neutron port mapping	12
2.9. OpenStack Architecture	13

Overview

[OpenStack](#) is a popular open source Infrastructure as a service project, covering compute, storage and network management. OpenStack can use OpenDaylight as its network management provider through the Modular Layer 2 (ML2) north-bound plug-in. OpenDaylight manages the network flows for the OpenStack compute nodes via the OVSDB south-bound plug-in. This page describes how to set that up, and how to tell when everything is working.

1. Installing OpenStack

Installing OpenStack is out of scope for this document, but to get started, it is useful to have a minimal multi-node OpenStack deployment.

The reference deployment we will use for this document is a 3 node cluster:

- One control node containing all of the management services for OpenStack (Nova, Neutron, Glance, Swift, Cinder, Keystone)
- Two compute nodes running nova-compute
- Neutron using the OVS back-end and vxlan for tunnels

Once you have installed OpenStack, verify that it is working by connecting to Horizon and performing a few operations. To check the Neutron configuration, create two instances on a private subnet bridging to your public network, and verify that you can connect to them, and that they can see each other.

2. Installing OpenDaylight

Table of Contents

OpenStack with OVSDB	2
OpenStack with GroupBasedPolicy	6
OpenStack with Virtual Tenant Network	12

OpenStack with OVSDB

Prerequisites: OpenDaylight requires Java 1.7.0.


- On the control host, [Download the latest OpenDaylight release](#) (at the time of writing, this is Beryllium)
- Uncompress it as root, and start OpenDaylight (you can start OpenDaylight by running karaf directly, but exiting from the shell will shut it down):

```
$ tar xvfz distribution-karaf-0.2.1-Helium-SR1.1.tar.gz
$ cd distribution-karaf-0.2.0-Helium
$ ./bin/start # Start OpenDaylight as a server process
```

- Connect to the Karaf shell, and install the odl-ovsdb-openstack bundle, dlux and their dependencies:

```
$ ./bin/client # Connect to OpenDaylight with the client
opendaylight-user@root> feature:install odl-base-all odl-aaa-authn odl-
restconf odl-nsf-all odl-adsal-northbound odl-mdsal-apidocs \
odl-ovsdb-openstack odl-ovsdb-northbound odl-dlux-core
```

- If everything is installed correctly, you should now be able to log in to the dlux interface on `http://$CONTROL_HOST:8181/dlux/index.html` - the default username and password is "admin/admin" (see screenshot below)



Ensuring OpenStack network state is clean

When using OpenDaylight as the Neutron back-end, ODL expects to be the only source of truth for Open vSwitch configuration. Because of this, it is necessary to remove existing OpenStack and Open vSwitch configurations to give OpenDaylight a clean slate.

- Delete instances

```
$ nova list
$ nova delete <instance names>
```

- Remove link from subnets to routers

```
$ neutron subnet-list
$ neutron router-list
$ neutron router-port-list <router name>
$ neutron router-interface-delete <router name> <subnet ID or name>
```

- Delete subnets, nets, routers

```
$ neutron subnet-delete <subnet name>
$ neutron net-list
$ neutron net-delete <net name>
$ neutron router-delete <router name>
```

- Check that all ports have been cleared - at this point, this should be an empty list

```
$ neutron port-list
```

Ensure Neutron is stopped

While Neutron is managing the OVS instances on compute and control nodes, OpenDaylight and Neutron can be in conflict. To prevent issues, we turn off Neutron server on the network controller, and Neutron's Open vSwitch agents on all hosts.

- Turn off neutron-server on control node

```
# systemctl stop neutron-server
```

- On each node in the cluster, shut down and disable Neutron's agent services to ensure that they do not restart after a reboot:

```
# systemctl stop neutron-openvswitch-agent
# systemctl disable neutron-openvswitch-agent
```

Configuring Open vSwitch to be managed by OpenDaylight

On each host (both compute and control nodes) we will clear the pre-existing Open vSwitch config and set OpenDaylight to manage the switch:

- Stop the Open vSwitch service, and clear existing OVSDB (ODL expects to manage vSwitches completely)

```
# systemctl stop openvswitch
# rm -rf /var/log/openvswitch/*
# rm -rf /etc/openvswitch/conf.db
# systemctl start openvswitch
```


- At this stage, your Open vSwitch configuration should be empty:

```
[root@dnearny-odl-compute2 ~]# ovs-vsctl show
9f3b38cb-eefc-4bc7-828b-084b1f66fbfd
    ovs_version: "2.1.3"
```

- Set OpenDaylight as the manager on all nodes

```
# ovs-vsctl set-manager tcp:${CONTROL_HOST}:6640
```

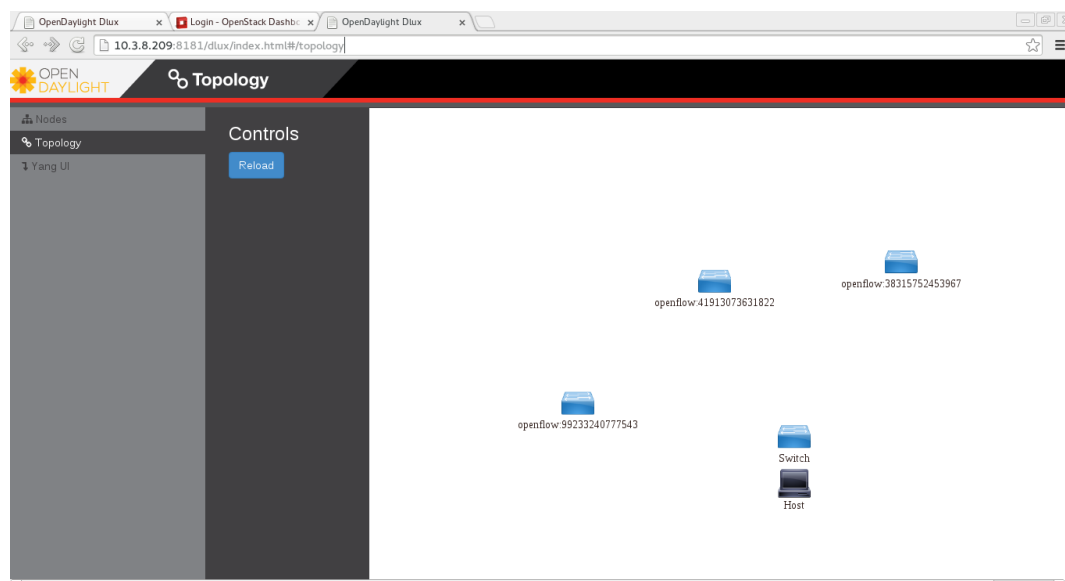
- You should now see a new section in your Open vSwitch configuration showing that you are connected to the OpenDaylight server, and OpenDaylight will automatically create a br-int bridge:

```
[root@dnearny-odl-compute2 ~]# ovs-vsctl show
9f3b38cb-eefc-4bc7-828b-084b1f66fbfd
    Manager "tcp:172.16.21.56:6640"
        is_connected: true
    Bridge br-int
        Controller "tcp:172.16.21.56:6633"
        fail_mode: secure
        Port br-int
            Interface br-int
    ovs_version: "2.1.3"
```

- (BUG WORKAROUND) If SELinux is enabled, you may not have a security context in place which allows Open vSwitch remote administration. If you do not see the result above (specifically, if you do not see "is_connected: true" in the Manager section), set SELinux to Permissive mode on all nodes and ensure it stays that way after boot:

```
# setenforce 0
# sed -i -e 's/SELINUX=enforcing/SELINUX=permissive/g' /etc/selinux/config
```

- Make sure all nodes, including the control node, are connected to OpenDaylight
- If you reload DLUX, you should now see that all of your Open vSwitch nodes are now connected to OpenDaylight



- If something has gone wrong, check `data/log/karaf.log` under the OpenDaylight distribution directory. If you do not see any interesting log entries, set logging for OVSDB to TRACE level inside Karaf and try again:

```
log:set TRACE ovsdb
```

Configuring Neutron to use OpenDaylight

Once you have configured the vSwitches to connect to OpenDaylight, you can now ensure that OpenStack Neutron is using OpenDaylight.

First, ensure that port 8080 (which will be used by OpenDaylight to listen for REST calls) is available. By default, swift-proxy-service listens on the same port, and you may need to move it (to another port or another host), or disable that service. I moved it to port 8081 by editing `/etc/swift/proxy-server.conf` and `/etc/cinder/cinder.conf`, modifying iptables appropriately, and restarting swift-proxy-service and OpenDaylight.

- Configure Neutron to use OpenDaylight's ML2 driver:

```
crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 mechanism_drivers
  opendaylight
crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 tenant_network_types
  vxlan

cat <<EOT>> /etc/neutron/plugins/ml2/ml2_conf.ini
[m2_odl]
password = admin
username = admin
url = http://${CONTROL_HOST}:8080/controller/nb/v2/neutron
EOT
```

- Reset Neutron's ML2 database

```
mysql -e "drop database if exists neutron_ml2;"
mysql -e "create database neutron_ml2 character set utf8;"
mysql -e "grant all on neutron_ml2.* to 'neutron'@'%';"
neutron-db-manage --config-file /usr/share/neutron/neutron-dist.conf --
config-file /etc/neutron/neutron.conf \
--config-file /etc/neutron/plugin.ini upgrade head
```

- Restart neutron-server:

```
systemctl start neutron-server
```

Verifying it works

- Verify that OpenDaylight's ML2 interface is working:

```
curl -u admin:admin http://${CONTROL_HOST}:8080/controller/nb/v2/neutron/
networks

{
  "networks" : [ ]
}
```

If this does not work or gives an error, check Neutron's log file in `<code>/var/log/neutron/server.log</code>`. Error messages here should give some clue as to what the problem is in the connection with OpenDaylight

- Create a net, subnet, router, connect ports, and start an instance using the Neutron CLI:

```
neutron router-create router1
neutron net-create private
neutron subnet-create private --name=private_subnet 10.10.5.0/24
neutron router-interface-add router1 private_subnet
nova boot --flavor <flavor> --image <image id> --nic net-id=<network id>
test1
nova boot --flavor <flavor> --image <image id> --nic net-id=<network id>
test2
```

At this point, you have confirmed that OpenDaylight is creating network end-points for instances on your network and managing traffic to them.

Congratulations! You're done!

OpenStack with GroupBasedPolicy

This section is for Application Developers and Network Administrators who are looking to integrate Group Based Policy with OpenStack.

To enable the **GBP** Neutron Mapper feature, at the karaf console:

```
feature:install odl-groupbasedpolicy-neutronmapper
```

Neutron Mapper has the following dependencies that are automatically loaded:

```
odl-neutron-service
```

Neutron Northbound implementing REST API used by OpenStack

```
odl-groupbasedpolicy-base
```

Base **GBP** feature set, such as policy resolution, data model etc.

```
odl-groupbasedpolicy-ofoverlay
```

For Lithium, **GBP** has one renderer, hence this is loaded by default.

REST calls from OpenStack Neutron are by the Neutron NorthBound project.

GBP provides the implementation of the [Neutron V2.0 API](#).

Features

List of supported Neutron entities:


- Port
- Network
 - Standard Internal

- External provider L2/L3 network
- Subnet
- Security-groups
- Routers
 - Distributed functionality with local routing per compute
 - External gateway access per compute node (dedicated port required)
 - Multiple routers per tenant
- FloatingIP NAT
- IPv4/IPv6 support

The mapping of Neutron entities to **GBP** entities is as follows:

Neutron Port

Figure 2.1. Neutron Port



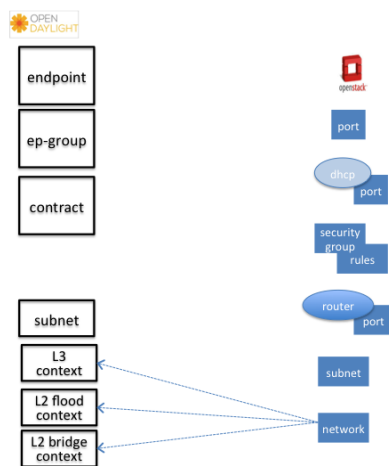
The Neutron port is mapped to an endpoint.

The current implementation supports one IP address per Neutron port.

An endpoint and L3-endpoint belong to multiple EndpointGroups if the Neutron port is in multiple Neutron Security Groups.

The key for endpoint is L2-bridge-domain obtained as the parent of L2-flood-domain representing Neutron network. The MAC address is from the Neutron port. An L3-endpoint is created based on L3-context (the parent of the L2-bridge-domain) and IP address of Neutron Port.

Neutron Network

Figure 2.2. Neutron Network

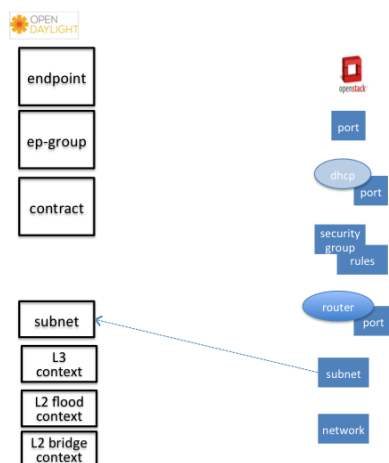
A Neutron network has the following characteristics:

- defines a broadcast domain
- defines a L2 transmission domain
- defines a L2 name space.

To represent this, a Neutron Network is mapped to multiple **GBP** entities. The first mapping is to an L2 flood-domain to reflect that the Neutron network is one flooding or broadcast domain. An L2-bridge-domain is then associated as the parent of L2 flood-domain. This reflects both the L2 transmission domain as well as the L2 addressing namespace.

The third mapping is to L3-context, which represents the distinct L3 address space. The L3-context is the parent of L2-bridge-domain.


Neutron Subnet

Figure 2.3. Neutron Subnet

Neutron subnet is associated with a Neutron network. The Neutron subnet is mapped to a **GBP** subnet where the parent of the subnet is L2-flood-domain representing the Neutron network.

Neutron Security Group

Figure 2.4. Neutron Security Group and Rules



GBP entity representing Neutron security-group is EndpointGroup.

Infrastructure EndpointGroups

Neutron-mapper automatically creates EndpointGroups to manage key infrastructure items such as:


- DHCP EndpointGroup - contains endpoints representing Neutron DHCP ports
- Router EndpointGroup - contains endpoints representing Neutron router interfaces
- External EndpointGroup - holds L3-endpoints representing Neutron router gateway ports, also associated with FloatingIP ports.

Neutron Security Group Rules

This mapping is most complicated among all others because Neutron security-group-rules are mapped to contracts with clauses, subjects, rules, action-refs, classifier-refs, etc. Contracts are used between endpoint groups representing Neutron Security Groups. For simplification it is important to note that Neutron security-group-rules are similar to a **GBP** rule containing:

- classifier with direction
- action of **allow**.

Neutron Routers

Figure 2.5. Neutron Router

Neutron router is represented as a L3-context. This treats a router as a Layer3 namespace, and hence every network attached to it a part of that Layer3 namespace.

This allows for multiple routers per tenant with complete isolation.

The mapping of the router to an endpoint represents the router's interface or gateway port.

The mapping to an EndpointGroup represents the internal infrastructure EndpointGroups created by the **GBP Neutron Mapper**

When a Neutron router interface is attached to a network/subnet, that network/subnet and its associated endpoints or Neutron Ports are seamlessly added to the namespace.

Neutron FloatingIP

When associated with a Neutron Port, this leverages the **GBP OfOverlay** renderer's NAT capabilities.

A dedicated *external* interface on each Nova compute host allows for distributed external access. Each Nova instance associated with a FloatingIP address can access the external network directly without having to route via the Neutron controller, or having to enable any form of Neutron distributed routing functionality.

Assuming the gateway provisioned in the Neutron Subnet command for the external network is reachable, the combination of **GBP Neutron Mapper** and **OfOverlay** renderer will automatically ARP for this default gateway, requiring no user intervention.

Troubleshooting within GBP


Logging level for the mapping functionality can be set for package `org.opendaylight.groupbasedpolicy.neutron.mapper`. An example of enabling TRACE logging level on karaf console:

```
log:set TRACE org.opendaylight.groupbasedpolicy.neutron.mapper
```

Neutron mapping example

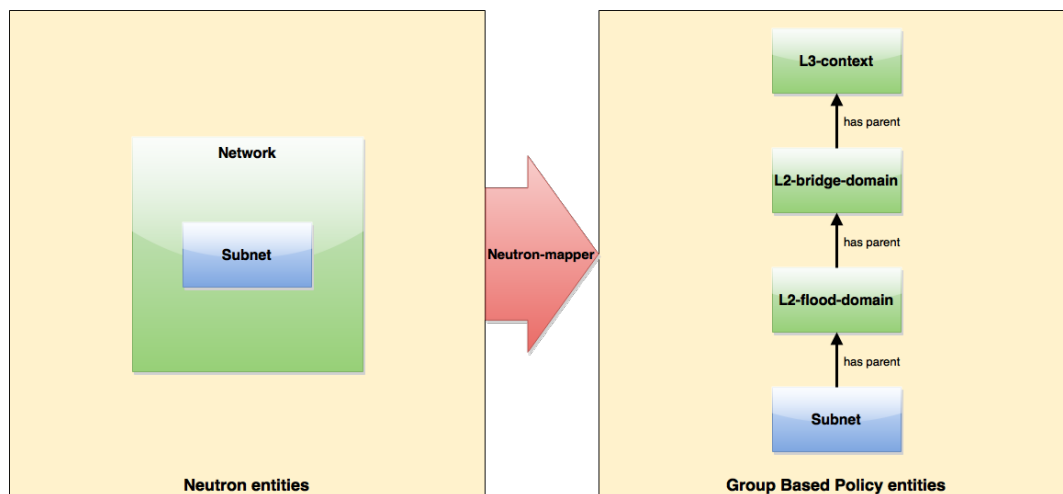
As an example for mapping can be used creation of Neutron network, subnet and port. When a Neutron network is created 3 **GBP** entities are created: l2-flood-domain, l2-bridge-domain, l3-context.

Figure 2.6. Neutron network mapping

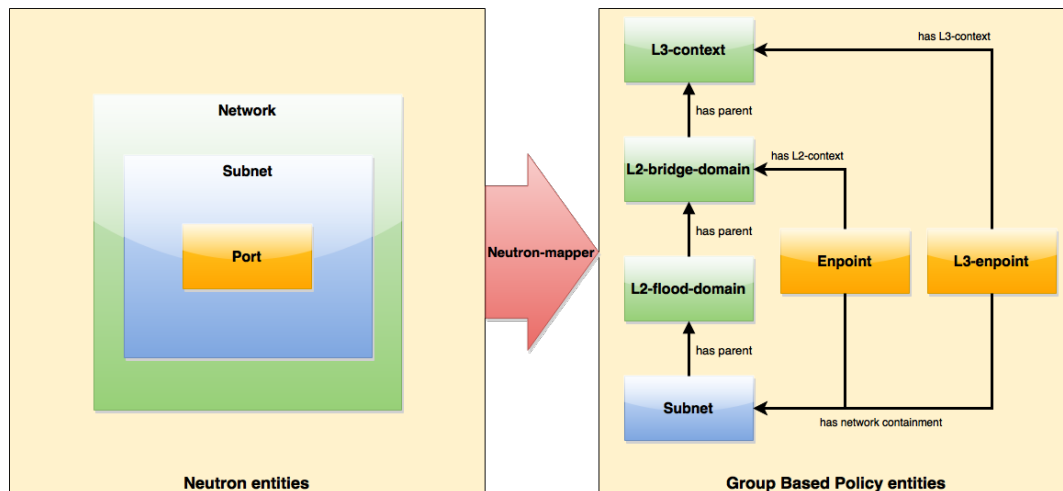


After an subnet is created in the network mapping looks like this.

Figure 2.7. Neutron subnet mapping



If an Neutron port is created in the subnet an endpoint and l3-endpoint are created. The endpoint has key composed from l2-bridge-domain and MAC address from Neutron port. A key of l3-endpoint is composed from l3-context and IP address. The network containment of endpoint and l3-endpoint points to the subnet.

Figure 2.8. Neutron port mapping

Configuring GBP Neutron

No intervention passed initial OpenStack setup is required by the user.

More information about configuration can be found in our DevStack demo environment on the [GBP wiki](#).

Administering or Managing GBP Neutron

For consistencies sake, all provisioning should be performed via the Neutron API. (CLI or Horizon).

The mapped policies can be augmented via the **GBP UX,UX**, to:

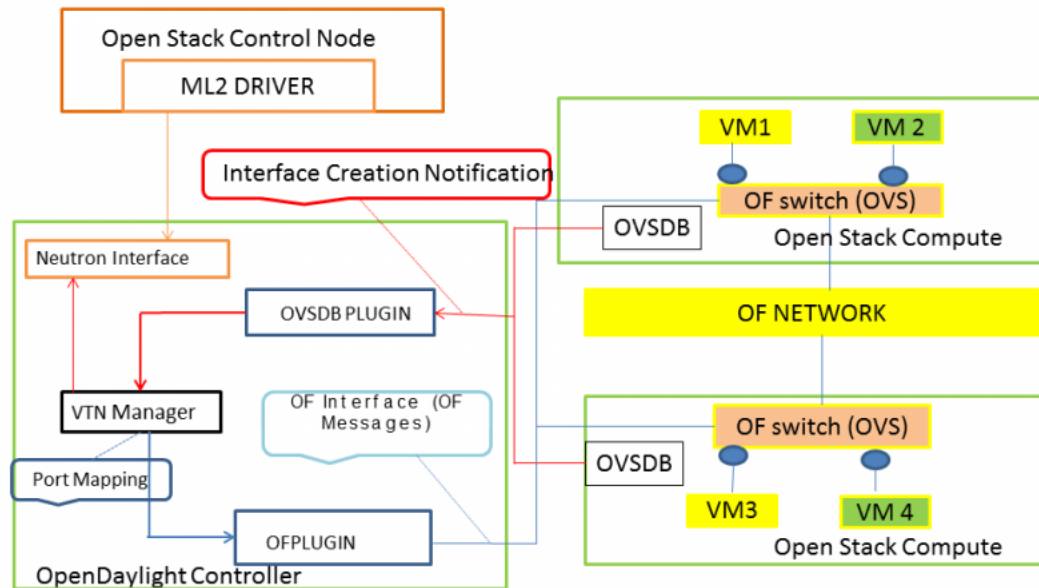
- Enable Service Function Chaining
- Add endpoints from outside of Neutron i.e. VMs/containers not provisioned in OpenStack
- Augment policies/contracts derived from Security Group Rules
- Overlay additional contracts or groupings

Tutorials

A DevStack demo environment can be found on the [GBP wiki](#).

OpenStack with Virtual Tenant Network

This section describes using OpenDaylight with the VTN manager feature providing network service for OpenStack. VTN manager utilizes the OVSDb southbound service and Neutron for this implementation. The below diagram depicts the communication of OpenDaylight and two virtual networks connected by an OpenFlow switch using this implementation.

Figure 2.9. OpenStack Architecture

Configure OpenStack to work with OpenDaylight(VTN Feature) using PackStack

Prerequisites to install OpenStack using PackStack

- Fresh CentOS 7.1 minimal install
- Use the below commands to disable and remove Network Manager in CentOS 7.1,


```
# systemctl stop NetworkManager
# systemctl disable NetworkManager
```
- To make SELINUX as permissive, please open the file "/etc/sysconfig/selinux" and change it as "SELINUX=permissive".
- After making selinux as permissive, please restart the CentOS 7.1 machine.

Steps to install OpenStack PackStack in CentOS 7.1

- To install OpenStack junos, use the following command,


```
# yum update -y
# yum -y install https://repos.fedorapeople.org/repos/openstack/openstack-juno/rdo-release-juno-1.noarch.rpm
```
- To install the packstack installer, please use the below command,

```
# yum -y install openstack-packstack
```

- To create all-in-one setup, please use the below command,

```
# packstack --allinone --provision-demo=n --provision-all-in-one-ovs-bridge=n
```

- This will end up with Horizon started successfully message.

Steps to install and deploy OpenDaylight in CentOS 7.1

- Download the latest lithium distribution code in the below link,

```
# wget https://nexus.opendaylight.org/content/groups/public/org/opendaylight/integration/distribution-karaf/0.3.1-Lithium-SR1/distribution-karaf-0.3.1-Lithium-SR1.zip
```

- Unzip the lithium distribution code by using the below command,

```
# unzip distribution-karaf-0.3.1-Lithium-SR1.zip
```

- Please do the below steps in the OpenDaylight to change jetty port,

- Change the jetty port from 8080 to something else as swift proxy of OpenStack is using it.
- Open the file "etc/jetty.xml" and change the jetty port from 8080 to 8910 (we have used 8910 as jetty port you can use any other number).
- Start VTN Manager and install odl-vtn-manager-neutron in it.
- Ensure all the required ports(6633/6653,6640 and 8910) are in the listen mode by using the command "netstat -tunpl" in OpenDaylight.

Steps to reconfigure OpenStack in CentOS 7.1

- Steps to stop Open vSwitch Agent and clean up ovs

```
# sudo systemctl stop neutron-openvswitch-agent
# sudo systemctl disable neutron-openvswitch-agent
# sudo systemctl stop openvswitch
# sudo rm -rf /var/log/openvswitch/*
# sudo rm -rf /etc/openvswitch/conf.db
# sudo systemctl start openvswitch
# sudo ovs-vsctl show
```

- Stop Neutron Server

```
# systemctl stop neutron-server
```

- Verify that OpenDaylight's ML2 interface is working:

```
curl -v admin:admin http://{CONTROL_HOST}:{PORT}/controller/nb/v2/neutron/networks

{
  "networks" : [ ]
}
```

```
}
```

If this does not work or gives an error, check Neutron's log file in `/var/log/neutron/server.log`. Error messages here should give some clue as to what the problem is in the connection with OpenDaylight

- Configure Neutron to use OpenDaylight's ML2 driver:

```
# sudo crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2
mechanism_drivers opendaylight
# sudo crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2
tenant_network_types local
# sudo crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 type_drivers
local
# sudo crudini --set /etc/neutron/dhcp_agent.ini DEFAULT ovs_use_veth True

# cat <<EOT | sudo tee -a /etc/neutron/plugins/ml2/ml2_conf.ini > /dev/null
[m2_odl]
password = admin
username = admin
url = http://{CONTROL_HOST}:{PORT}/controller/nb/v2/neutron
EOT
```

- Reset Neutron's ML2 database

```
# sudo mysql -e "drop database if exists neutron_ml2;"
# sudo mysql -e "create database neutron_ml2 character set utf8;"
# sudo mysql -e "grant all on neutron_ml2.* to 'neutron'@'%';"
# sudo neutron-db-manage --config-file /usr/share/neutron/neutron-dist.conf
--config-file /etc/neutron/neutron.conf --config-file /etc/neutron/
plugin.ini upgrade head
```

- Start Neutron Server

```
# sudo systemctl start neutron-server
```

- Restart the Neutron DHCP service

```
# system restart neutron-dhcp-agent.service
```

- At this stage, your Open vSwitch configuration should be empty:

```
[root@dneary-odl-compute2 ~]# ovs-vsctl show
686989e8-7113-4991-a066-1431e7277e1f
    ovs_version: "2.3.1"
```

- Set OpenDaylight as the manager on all nodes

```
# ovs-vsctl set-manager tcp:127.0.0.1:6640
```

- You should now see a section in your Open vSwitch configuration showing that you are connected to the OpenDaylight server, and OpenDaylight will automatically create a br-int bridge:

```
[root@dneary-odl-compute2 ~]# ovs-vsctl show
686989e8-7113-4991-a066-1431e7277e1f
    Manager "tcp:127.0.0.1:6640"
        is_connected: true
    Bridge br-int
        Controller "tcp:127.0.0.1:6633"
```

```
is_connected: true
fail_mode: secure
Port "ens33"
  Interface "ens33"
ovs_version: "2.3.1"
```

- Add the default flow to OVS to forward packets to controller when there is a table-miss,

```
ovs-ofctl --protocols=OpenFlow13 add-flow br-int
priority=0,actions=output:CONTROLLER
```

- Please see the [VTN OpenStack PackStack support guide on the wiki](#) to create VM's from OpenStack Horizon GUI.

Implementation details

VTN Manager:

Install **odl-vtn-manager-neutron** feature which provides the integration with Neutron interface.

```
feature:install odl-vtn-manager-neutron
```

It subscribes to the events from Open vSwitch and also implements the Neutron requests received by OpenDaylight.

Functional Behavior

StartUp:

- The ML2 implementation for OpenDaylight will ensure that when Open vSwitch is started, the ODL_IP_ADDRESS configured will be set as manager.
- When OpenDaylight receives the update of the Open vSwitch on port 6640 (manager port), VTN Manager handles the event and adds a bridge with required port mappings to the Open vSwitch at the OpenStack node.
- When Neutron starts up, a new network create is POSTed to OpenDaylight, for which VTN Manager creates a Virtual Tenant Network.
- **Network and Sub-Network Create:** Whenever a new sub network is created, VTN Manager will handle the same and create a vbridge under the VTN.
- **VM Creation in OpenStack:** The interface mentioned as integration bridge in the configuration file will be added with more interfaces on creation of a new VM in OpenStack and the network is provisioned for it by the VTN Neutron feature. The addition of a new port is captured by the VTN Manager and it creates a vbridge interface with port mapping for the particular port. When the VM starts to communicate with other VMs, the VTN Manger will install flows in the Open vSwitch and other OpenFlow switches to facilitate communication between them.



Note

To use this feature, VTN feature should be installed

Reference

https://wiki.opendaylight.org/images/5/5c/Integration_of_vtn_and_ovsdb_for_helium.pdf