
OpenStack-Ansible Documentation: **os_neutron role** *Release 18.1.0.dev477*

OpenStack-Ansible Contributors

Apr 04, 2022

CONTENTS

1	Configuring the Networking service (neutron) (optional)	1
1.1	Firewall service (optional)	1
1.2	Virtual private network service (optional)	3
1.3	BGP Dynamic Routing service (optional)	4
1.4	SR-IOV Support (optional)	4
2	Scenario - Using Open vSwitch	7
2.1	Overview	7
2.2	Recommended reading	7
2.3	Prerequisites	7
2.4	Configuring bridges (Linux Bridge)	8
2.5	Configuring bridges (Open vSwitch)	8
2.6	OpenStack-Ansible user variables	9
2.7	Open Virtual Switch (OVS) commands	11
2.8	Notes	13
3	Scenario - Using Open vSwitch w/ ASAP² (Direct Mode)	15
3.1	Overview	15
3.2	Recommended reading	15
3.3	Prerequisites	16
3.4	Deployment	16
3.5	Post-Deployment	17
3.6	Verify operation	18
4	Scenario - Using Open vSwitch with DVR	21
4.1	Overview	21
4.2	Recommended reading	21
4.3	Prerequisites	21
4.4	OpenStack-Ansible user variables	21
5	Scenario - Using Open vSwitch w/ DPDK	25
5.1	Overview	25
5.2	Recommended reading	25
5.3	Prerequisites	25
5.4	NUMA topology	26
5.5	Hugepage configuration	28
5.6	OpenStack-Ansible variables and overrides	29
5.7	Flavor configuration	31
5.8	OpenStack-Ansible user variables	31

5.9	Post-installation	32
6	Scenario - Using Open vSwitch with SFC	33
6.1	Overview	33
6.2	Recommended reading	33
6.3	Prerequisites	33
6.4	OpenStack-Ansible user variables	33
7	Scenario - Open Virtual Network (OVN)	35
7.1	Overview	35
7.2	Recommended reading	35
7.3	Prerequisites	36
7.4	OpenStack-Ansible user variables	36
7.5	(Optional) DVR or Distributed L3 routing	38
7.6	Open Virtual Network (OVN) commands	39
7.7	Notes	40
8	Scenario - Using the Nuage neutron plugin	43
8.1	Introduction	43
8.2	Prerequisites	43
8.3	Configure Nuage neutron plugin	43
8.4	Installation	45
9	Scenario - VMware NSX Plugin	47
9.1	Introduction	47
9.2	Prerequisites	47
9.3	Configure Neutron to use the NSX plugin	47
9.4	Installation	48
10	Scenario - Using the networking-calico Neutron plugin	49
10.1	Introduction	49
10.2	Prerequisites	49
10.3	Configure OSA Environment for Project Calico	49
10.4	Configure networking-calico Neutron Plugin	50
10.5	Installation	50
11	Scenario - OpenDaylight and Open vSwitch	51
11.1	Overview	51
11.2	Recommended reading	51
11.3	Prerequisites	51
11.4	OpenStack-Ansible user variables	51
11.5	L3 configuration	52
11.6	SFC configuration	53
11.7	BGPVPN configuration	53
11.8	Security information	53
12	Scenario - Networking Generic Switch	55
12.1	Overview	55
12.2	Recommended reading	55
12.3	Prerequisites	55
12.4	OpenStack-Ansible user variables	55
12.5	Notes	56

13	Default variables	57
14	Dependencies	67
15	Example playbook	69
16	Tags	71

CONFIGURING THE NETWORKING SERVICE (NEUTRON) (OPTIONAL)

The OpenStack Networking service (neutron) includes the following services:

Firewall as a Service (FWaaS) Provides a software-based firewall that filters traffic from the router.

VPN as a Service (VPNaaS) Provides a method for extending a private network across a public network.

BGP Dynamic Routing service Provides a means for advertising self-service (private) network prefixes to physical network devices that support BGP.

SR-IOV Support Provides the ability to provision virtual or physical functions to guest instances using SR-IOV and PCI passthrough. (Requires compatible NICs)

1.1 Firewall service (optional)

The following procedure describes how to modify the `/etc/openstack_deploy/user_variables.yml` file to enable FWaaS.

1.1.1 Deploying FWaaS v1

Note: The FWaaS v1 API is deprecated upstream. While FWaaS v1.0 is still maintained, new features will be implemented in FWaaS v2.0 API.

1. Override the default list of neutron plugins to include `firewall`:

```
neutron_plugin_base:  
- firewall  
- ...
```

2. `neutron_plugin_base` is as follows:

```
neutron_plugin_base:  
- router  
- firewall  
- vpnaas  
- metering  
- qos
```

3. Execute the neutron install playbook in order to update the configuration:

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-neutron-install.yml
```

4. Execute the horizon install playbook to show the FWaaS panels:

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-horizon-install.yml
```

The FWaaS default configuration options may be changed through the [conf override](#) mechanism using the `neutron_neutron_conf_overrides` dict.

1.1.2 Deploying FWaaS v2

FWaaS v2 is the next generation Neutron firewall service and will provide a rich set of APIs for securing OpenStack networks. It is still under active development.

Refer to the [FWaaS 2.0 API specification](#) for more information on these FWaaS v2 features.

FWaaS v2 requires the use of Open vSwitch. To deploy an environment using Open vSwitch for virtual networking, please refer to the following documentation:

- [Scenario - Using Open vSwitch](#)
- [Scenario - Using Open vSwitch with DVR](#)

Follow the steps below to deploy FWaaS v2:

Note: FWaaS v1 and v2 cannot be deployed simultaneously.

1. Add the FWaaS v2 plugin to the `neutron_plugin_base` variable in `/etc/openstack_deploy/user_variables.yml`:

```
neutron_plugin_base:
  - router
  - metering
  - firewall_v2
```

Ensure that `neutron_plugin_base` includes all of the plugins that you want to deploy with neutron in addition to the `firewall_v2` plugin.

2. Run the neutron playbook to deploy the FWaaS v2 service plugin

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-neutron-install.yml
```


1.2 Virtual private network service (optional)

The following procedure describes how to modify the `/etc/openstack_deploy/user_variables.yml` file to enable VPNaaS.

1. Override the default list of neutron plugins to include vpnaas:

```
neutron_plugin_base:
  - router
  - metering
```

2. neutron_plugin_base is as follows:

```
neutron_plugin_base:
  - router
  - metering
  - vpnaas
```

3. Override the default list of specific kernel modules in order to include the necessary modules to run ipsec:

```
openstack_host_specific_kernel_modules:
  - { name: "ebtables", pattern: "CONFIG_BRIDGE_NF_EBTABLES=", ↵
↵group: "network_hosts" }
  - { name: "af_key", pattern: "CONFIG_NET_KEY=", group: "network_
↵hosts" }
  - { name: "ah4", pattern: "CONFIG_INET_AH=", group: "network_hosts
↵" }
  - { name: "ipcomp", pattern: "CONFIG_INET_IPCOMP=", group:
↵"network_hosts" }
```

4. Execute the openstack hosts setup in order to load the kernel modules at boot and runtime in the network hosts

```
# openstack-ansible openstack-hosts-setup.yml --limit network_hosts\
--tags "openstack_hosts-config"
```

5. Execute the neutron install playbook in order to update the configuration:

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-neutron-install.yml
```

6. Execute the horizon install playbook to show the VPNaaS panels:

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-horizon-install.yml
```

The VPNaaS default configuration options are changed through the `conf override` mechanism using the `neutron_neutron_conf_overrides` dict.

1.3 BGP Dynamic Routing service (optional)

The [BGP Dynamic Routing](#) plugin for neutron provides BGP speakers which can advertise OpenStack project network prefixes to external network devices, such as routers. This is especially useful when coupled with the [subnet pools](#) feature, which enables neutron to be configured in such a way as to allow users to create self-service [segmented IPv6 subnets](#).

The following procedure describes how to modify the `/etc/openstack_deploy/user_variables.yml` file to enable the BGP Dynamic Routing plugin.

1. Add the BGP plugin to the `neutron_plugin_base` variable in `/etc/openstack_deploy/user_variables.yml`:

```
neutron_plugin_base:
- ...
- neutron_dynamic_routing.services.bgp.bgp_plugin.BgpPlugin
```

Ensure that `neutron_plugin_base` includes all of the plugins that you want to deploy with neutron in addition to the BGP plugin.

2. Execute the neutron install playbook in order to update the configuration:

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-neutron-install.yml
```

1.4 SR-IOV Support (optional)

The following procedure describes how to modify the OpenStack-Ansible configuration to enable Neutron SR-IOV support.

1. Define SR-IOV capable physical host interface for a provider network

As part of every Openstack-Ansible installation, all provider networks known to Neutron need to be configured inside the `/etc/openstack_deploy/openstack_user_config.yml` file. For each supported network type (e.g. vlan), the attribute `sriov_host_interfaces` can be defined to map ML2 network names (`net_name` attribute) to one or many physical interfaces. Additionally, the network will need to be assigned to the `neutron_sriov_nic_agent` container group.

Example configuration:

```
provider_networks
- network:
  container_bridge: "br-vlan"
  container_type: "veth"
  container_interface: "eth11"
  type: "vlan"
  range: "1000:2000"
  net_name: "physnet1"
  sriov_host_interfaces: "p1p1,p4p1"
  group_binds:
    - neutron_linuxbridge_agent
    - neutron_sriov_nic_agent
```

2. Configure Nova

With SR-IOV, Nova uses PCI passthrough to allocate VFs and PFs to guest instances. Virtual Functions (VFs) represent a slice of a physical NIC, and are passed as virtual NICs to guest instances. Physical Functions (PFs), on the other hand, represent an entire physical interface and are passed through to a single guest.

To use PCI passthrough in Nova, the `PciPassthroughFilter` filter needs to be added to the `conf override nova_scheduler_default_filters`. Finally, PCI devices available for passthrough need to be allow via the `conf override nova_pci_passthrough_whitelist`.

Possible options which can be configured:

```
# Single device configuration
nova_pci_passthrough_whitelist: '{ "physical_network": "physnet1",
  ↪ "devname": "plp1" }'
```

```
# Multi device configuration
nova_pci_passthrough_whitelist: '[{"physical_network": "physnet1",
  ↪ "devname": "plp1"}, {"physical_network": "physnet1", "devname": "p4p1"}
  ↪]'
```

```
# Whitelisting by PCI Device Location
# The example pattern for the bus location '0000:04:*. *' is very wide.
  ↪ Make sure that
# no other, unintended devices, are whitelisted (see lspci -nn)
nova_pci_passthrough_whitelist: '{"address": "0000:04:*. *", "physical_
  ↪ network": "physnet1"}'
```

```
# Whitelisting by PCI Device Vendor
# The example pattern limits matches to PCI cards with vendor id 8086_
  ↪ (Intel) and
# product id 10ed (82599 Virtual Function)
nova_pci_passthrough_whitelist: '{"vendor_id": "8086", "product_id":
  ↪ "10ed", "physical_network": "physnet1"}'
```

```
# Additionally, devices can be matched by their type, VF or PF, using_
  ↪ the dev_type parameter
# and type-VF or type-PF options
nova_pci_passthrough_whitelist: '{"vendor_id": "8086", "product_id":
  ↪ "10ed", "dev_type": "type-VF", "physical_network": "physnet1"}'
```

It is recommended to use whitelisting by either the Linux device name (`devname` attribute) or by the PCI vendor and product id combination (`vendor_id` and `product_id` attributes)

3. Enable the SR-IOV ML2 plugin

The `conf override neutron_plugin_type` variable defines the core ML2 plugin, and only one plugin can be defined at any given time. The `conf override neutron_plugin_types` variable can contain a list of additional ML2 plugins to load. Make sure that only compatible ML2 plugins are loaded at all times. The SR-IOV ML2 plugin is known to work with the linuxbridge (`m12.lxb`) and openvswitch (`m12.ovs`) ML2 plugins. `m12.lxb` is the standard activated core ML2 plugin.

```
neutron_plugin_types:
  - m12.sriov
```

4. Execute the Neutron install playbook in order to update the configuration:

```
# cd /opt/openstack-ansible/playbooks
# openstack-ansible os-neutron-install.yml
# openstack-ansible os-nova-install.yml
```

5. Check Neutron SR-IOV agent state

After the playbooks have finished configuring Neutron and Nova, the new Neutron Agent state can be verified with:

```
# neutron agent-list --agent_type 'NIC Switch agent'
+-----+-----+-----+-----+-----+-----+
↪+-----+-----+-----+-----+-----+-----+
| id | agent_type | host |
↪| alive | admin_state_up | binary |
+-----+-----+-----+-----+-----+-----+
↪+-----+-----+-----+-----+-----+-----+
| 3012ff0e-de35-447b-aff6-fdb55b04c518 | NIC Switch agent | compute01 |
↪| :- ) | True | neutron-sriov-nic-agent |
| bb0c0385-394d-4e72-8bfe-26fd020df639 | NIC Switch agent | compute02 |
↪| :- ) | True | neutron-sriov-nic-agent |
+-----+-----+-----+-----+-----+-----+
↪+-----+-----+-----+-----+-----+-----+
```

Deployers can make changes to the SR-IOV nic agent default configuration options via the `neutron_sriov_nic_agent_ini_overrides` dict. Review the documentation on the [conf override](#) mechanism for more details.

SCENARIO - USING OPEN VSWITCH

2.1 Overview

Operators can choose to utilize Open vSwitch instead of Linux Bridges for the neutron ML2 agent. This offers different capabilities and integration points with neutron. This document outlines how to set it up in your environment.

2.2 Recommended reading

We recommend that you read the following documents before proceeding:

- Neutron documentation on Open vSwitch OpenStack deployments: <https://docs.openstack.org/liberty/networking-guide/scenario-classic-ovs.html>
- Blog post on how OpenStack-Ansible works with Open vSwitch: <https://medium.com/@travistruman/configuring-openstack-ansible-for-open-vswitch-b7e70e26009d>

2.3 Prerequisites

All compute nodes must have bridges configured:

- `br-mgmt`
- `br-vlan` (optional - used for vlan networks)
- `br-vxlan` (optional - used for vxlan tenant networks)
- `br-storage` (optional - used for certain storage devices)

For more information see: <https://docs.openstack.org/project-deploy-guide/openstack-ansible/newton/targethosts-networkconfig.html>

These bridges may be configured as either a Linux Bridge (which would connect to the Open vSwitch controlled by neutron) or as an Open vSwitch.

2.4 Configuring bridges (Linux Bridge)

The following is an example of how to configure a bridge (example: br-mgmt) with a Linux Bridge on Ubuntu 16.04 LTS:

```
/etc/network/interfaces
```

```
auto lo
iface lo inet loopback

# Management network
auto eth0
iface eth0 inet manual

# VLAN network
auto eth1
iface eth1 inet manual

source /etc/network/interfaces.d/*.cfg
```

```
/etc/network/interfaces.d/br-mgmt.cfg
```

```
# OpenStack Management network bridge
auto br-mgmt
iface br-mgmt inet static
    bridge_stp off
    bridge_waitport 0
    bridge_fd 0
    bridge_ports eth0
    address MANAGEMENT_NETWORK_IP
    netmask 255.255.255.0
```

One br-<type>.cfg is required for each bridge. VLAN interfaces can be used to back the br-<type> bridges if there are limited physical adapters on the system.

2.5 Configuring bridges (Open vSwitch)

Another configuration method routes everything with Open vSwitch. The bridge (example: br-mgmt) can be an Open vSwitch itself.

The following is an example of how to configure a bridge (example: br-mgmt) with Open vSwitch on Ubuntu 16.04 LTS: *

```
/etc/network/interfaces
```

```
auto lo
iface lo inet loopback

source /etc/network/interfaces.d/*.cfg

# Management network
allow-br-mgmt eth0
iface eth0 inet manual
    ovs_bridge br-mgmt
    ovs_type OVSPort
```

(continues on next page)

(continued from previous page)

```
# VLAN network
allow-br-vlan eth1
iface eth1 inet manual
    ovs_bridge br-vlan
    ovs_type OVSPort
```

/etc/network/interfaces.d/br-mgmt.cfg

```
# OpenStack Management network bridge
auto br-mgmt
allow-ovs br-mgmt
iface br-mgmt inet static
    address MANAGEMENT_NETWORK_IP
    netmask 255.255.255.0
    ovs_type OVSBridge
    ovs_ports eth0
```

One `br-<type>.cfg` is required for each bridge. VLAN interfaces can be used to back the `br-<type>` bridges if there are limited physical adapters on the system.

Warning: There is a bug in Ubuntu 16.04 LTS where the Open vSwitch service won't start properly when using systemd. The bug and workaround are discussed here: <http://www.opencloudblog.com/?p=240>

2.6 OpenStack-Ansible user variables

Create a group var file for your network hosts `/etc/openstack_deploy/group_vars/network_hosts`. It has to include:

```
# Ensure the openvswitch kernel module is loaded
openstack_host_specific_kernel_modules:
  - name: "openvswitch"
    pattern: "CONFIG_OPENVSWITCH"
```

Specify provider network definitions in your `/etc/openstack_deploy/openstack_user_config.yml` that define one or more Neutron provider bridges and related configuration:

Note: Bridges specified here will be created automatically. If `network_interface` is defined, the interface will be placed into the bridge automatically.

```
- network:
    container_bridge: "br-provider"
    container_type: "veth"
    type: "vlan"
    range: "101:200,301:400"
    net_name: "physnet1"
    network_interface: "bond1"
    group_binds:
      - neutron_openvswitch_agent
- network:
```

(continues on next page)

(continued from previous page)

```
container_bridge: "br-provider2"
container_type: "veth"
type: "vlan"
range: "203:203,467:500"
net_name: "physnet2"
network_interface: "bond2"
group_binds:
  - neutron_openvswitch_agent
```

When using flat provider networks, modify the network type accordingly:

```
- network:
  container_bridge: "br-publicnet"
  container_type: "veth"
  type: "flat"
  net_name: "flat"
  group_binds:
    - neutron_openvswitch_agent
```

Specify an overlay network definition in your `/etc/openstack_deploy/openstack_user_config.yml` that defines overlay network-related configuration:

Note: The bridge name should correspond to a pre-created Linux bridge or OVS bridge.

```
- network:
  container_bridge: "br-vxlan"
  container_type: "veth"
  container_interface: "eth10"
  ip_from_q: "tunnel"
  type: "vxlan"
  range: "1:1000"
  net_name: "vxlan"
  group_binds:
    - neutron_openvswitch_agent
```

Set the following user variables in your `/etc/openstack_deploy/user_variables.yml`:

```
neutron_plugin_type: ml2.ovs
neutron_ml2_drivers_type: "flat,vlan,vxlan"
```

The overrides are instructing Ansible to deploy the OVS mechanism driver and associated OVS components. This is done by setting `neutron_plugin_type` to `ml2.ovs`.

The `neutron_ml2_drivers_type` override provides support for all common type drivers supported by OVS.

If provider network overrides are needed on a global or per-host basis, the following format can be used in `user_variables.yml` or per-host in `openstack_user_config.yml`.

Note: These overrides are not normally required when defining global provider networks in the `openstack_user_config.yml` file.

```

# When configuring Neutron to support vxlan tenant networks and
# vlan provider networks the configuration may resemble the following:
neutron_provider_networks:
  network_types: "vxlan"
  network_vxlan_ranges: "1:1000"
  network_vlan_ranges: "physnet1:102:199"
  network_mappings: "physnet1:br-provider"
  network_interface_mappings: "br-provider:bond1"

# When configuring Neutron to support only vlan tenant networks and
# vlan provider networks the configuration may resemble the following:
neutron_provider_networks:
  network_types: "vlan"
  network_vlan_ranges: "physnet1:102:199"
  network_mappings: "physnet1:br-provider"
  network_interface_mappings: "br-provider:bond1"

# When configuring Neutron to support multiple vlan provider networks
# the configuration may resemble the following:
neutron_provider_networks:
  network_types: "vlan"
  network_vlan_ranges: "physnet1:102:199,physnet2:2000:2999"
  network_mappings: "physnet1:br-provider,physnet2:br-provider2"
  network_interface_mappings: "br-provider:bond1,br-provider2:bond2"

# When configuring Neutron to support multiple vlan and flat provider
# networks the configuration may resemble the following:
neutron_provider_networks:
  network_flat_networks: "*"
  network_types: "vlan"
  network_vlan_ranges: "physnet1:102:199,physnet2:2000:2999"
  network_mappings: "physnet1:br-provider,physnet2:br-provider2"
  network_interface_mappings: "br-provider:bond1,br-provider2:bond2"

```

2.7 Open Virtual Switch (OVS) commands

The following commands can be used to provide useful information about the state of Open vSwitch networking and configurations.

The `ovs-vsctl show` command provides information about the virtual switches and connected ports currently configured on the host:

```

root@infra01:~# ovs-vsctl show
4ef304ff-b803-4d09-95f5-59a076323949
  Manager "ptcp:6640:127.0.0.1"
    is_connected: true
  Bridge br-int
    Controller "tcp:127.0.0.1:6633"
      is_connected: true
    fail_mode: secure
    Port "tap2e7e0507-e4"
      tag: 2
      Interface "tap2e7e0507-e4"
        type: internal

```

(continues on next page)

(continued from previous page)

```
Port int-br-vlan
  Interface int-br-vlan
    type: patch
    options: {peer=phy-br-provider}
Port br-int
  Interface br-int
    type: internal
Port "tap7796ab3d-e9"
  tag: 5
  Interface "tap7796ab3d-e9"
    type: internal
Port patch-tun
  Interface patch-tun
    type: patch
    options: {peer=patch-int}
Bridge br-tun
  Controller "tcp:127.0.0.1:6633"
    is_connected: true
  fail_mode: secure
  Port "vxlan-acldf015"
    Interface "vxlan-acldf015"
      type: vxlan
      options: {df_default="true", in_key=flow, local_ip="172.29.
↪240.20", out_key=flow, remote_ip="172.29.240.21"}
  Port patch-int
    Interface patch-int
      type: patch
      options: {peer=patch-tun}
  Port "vxlan-acldf017"
    Interface "vxlan-acldf017"
      type: vxlan
      options: {df_default="true", in_key=flow, local_ip="172.29.
↪240.20", out_key=flow, remote_ip="172.29.240.23"}
  Port br-tun
    Interface br-tun
      type: internal
Bridge br-provider
  Controller "tcp:127.0.0.1:6633"
    is_connected: true
  fail_mode: secure
  Port "ens192"
    Interface "ens192"
  Port br-provider
    Interface br-provider
      type: internal
  Port phy-br-provider
    Interface phy-br-provider
      type: patch
      options: {peer=int-br-provider}
ovs_version: "2.10.0"
```

Additional commands can be found in upstream Open vSwitch documentation.

2.8 Notes

The `neutron-openvswitch-agent` service will check in as an agent and can be observed using the `openstack network agent list` command:

```
root@infra01-utility-container-ce1509fd:~# openstack network agent list --
↪agent-type open-vswitch
+-----+-----+-----+-----+
↪+-----+-----+-----+-----+
| ID | Agent Type | Host |
↪| Availability Zone | Alive | State | Binary |
+-----+-----+-----+-----+
↪+-----+-----+-----+-----+
| 4dcef710-ec0c-4925-a940-dc319cd6849f | Open vSwitch agent | compute03 |
↪| None | :- ) | UP | neutron-openvswitch-agent |
| 5e1f8670-b90e-49c3-84ff-e981aeccb171 | Open vSwitch agent | compute02 |
↪| None | :- ) | UP | neutron-openvswitch-agent |
| 78746672-d77a-4d8a-bb48-f659251fa246 | Open vSwitch agent | compute01 |
↪| None | :- ) | UP | neutron-openvswitch-agent |
| eebab5da-3ef5-4582-84c5-f29e2472a44a | Open vSwitch agent | infra01 |
↪| None | :- ) | UP | neutron-openvswitch-agent |
+-----+-----+-----+-----+
↪+-----+-----+-----+-----+
```


SCENARIO - USING OPEN VSWITCH W/ ASAP² (DIRECT MODE)

3.1 Overview

With appropriate hardware, operators can choose to utilize ASAP²-accelerated Open vSwitch instead of unaccelerated Open vSwitch for the Neutron virtual network infrastructure. ASAP² technology offloads packet processing onto hardware built into the NIC rather than using the CPU of the host. It requires careful consideration and planning before implementing. This document outlines how to set it up in your environment.

Note: ASAP² is a proprietary feature provided with certain Mellanox NICs, including the ConnectX-4 Lx and ConnectX-5. Future support is not guaranteed. This feature is considered *EXPERIMENTAL* and should not be used for production workloads. There is no guarantee of upgradability or backwards compatibility.

Note: Hardware offloading is not compatible with the `openvswitch` firewall driver. To ensure flows are offloaded, port security must be disabled. Information on disabling port security is discussed later in this document.

3.2 Recommended reading

This guide is a variation of the standard Open vSwitch and SR-IOV deployment guides available at:

- https://docs.openstack.org/openstack-ansible-os_neutron/latest/app-openvswitch.html
- https://docs.openstack.org/openstack-ansible-os_neutron/latest/configure-network-services.html#sr-io-v-support-optional

The following resources may also be helpful:

- <https://docs.openstack.org/neutron/latest/admin/config-sriov.html>
- <https://docs.openstack.org/neutron/latest/admin/config-ovs-offload.html>
- https://www.mellanox.com/related-docs/prod_software/ASAP2_Hardware_Offloading_for_vSwitches_User_Manual_v4.4.pdf

3.3 Prerequisites

To enable SR-IOV and PCI passthrough capabilities on a Linux platform, ensure that VT-d/VT-x are enabled for Intel processors and AMD-V/AMD-Vi are enabled for AMD processors. Such features are typically enabled in the BIOS.

On an Intel platform, the following kernel parameters are required and can be added to the GRUB configuration:

```
GRUB_CMDLINE_LINUX="... iommu=pt intel_iommu=on"
```

On an AMD platform, use these parameters instead:

```
GRUB_CMDLINE_LINUX="... iommu=pt amd_iommu=on"
```

Update GRUB and reboot the host(s).

SR-IOV provides virtual functions (VFs) that can be presented to instances as network interfaces and are used in lieu of tuntap interfaces. Configuration of VFs is outside the scope of this guide. The following links may be helpful:

- <https://community.mellanox.com/s/article/getting-started-with-mellanox-asap-2>
- <https://community.mellanox.com/s/article/howto-configure-sr-iov-for-connectx-4-connectx-5-with-kvmethernet-x>

3.4 Deployment

Configure your networking according the Open vSwitch implementation docs:

- [Scenario - Using Open vSwitch](#)

Note: At this time, only a single (non-bonded) interface is supported.

An example provider network configuration has been provided below:

```
- network:
  container_bridge: "br-provider"
  container_type: "veth"
  type: "vlan"
  range: "700:709"
  net_name: "physnet1"
  network_interface: "ens4f0"
  group_binds:
    - neutron_openvswitch_agent
```

Add a nova_pci_passthrough_whitelist entry to user_variables.yml, where devname is the name of the interface connected to the provider bridge and physical_network is the name of the provider network.

```
nova_pci_passthrough_whitelist: '{"devname":"ens4f0","physical_network":
  ↳"physnet1"}'
```

Note: In the respective network block configured in `openstack_user_config.yml`, `devname` corresponds to `network_interface` and `physical_network` corresponds to `net_name`.

To enable the `openvswitch` firewall driver rather than the default `iptables_hybrid` firewall driver, add the following overrides to `user_variables.yml`:

```
neutron_ml2_conf_ini_overrides:
  securitygroup:
    firewall_driver: openvswitch
neutron_openvswitch_agent_ini_overrides:
  securitygroup:
    firewall_driver: openvswitch
```

Note: Hardware-offloaded flows are **not** activated for ports utilizing security groups or port security. Be sure to disable port security *and* security groups on individual ports or networks when hardware offloading is required.

Once the OpenStack cluster is configured, start the OpenStack deployment as listed in the OpenStack-Ansible Install guide by running all playbooks in sequence on the deployment host.

3.5 Post-Deployment

Once the deployment is complete, create the VFs that will be used for SR-IOV. In this example, the physical function (PF) is `ens4f0`. It will simultaneously be connected to the Neutron provider bridge `br-provider`.

1. On each compute node, determine the maximum number of VFs a PF can support:

```
# cat /sys/class/net/ens4f0/device/sriov_totalvfs
```

Note: To adjust `sriov_totalvfs` please refer to Mellanox documentation.

2. On each compute node, create the VFs:

```
# echo '8' > /sys/class/net/ens4f0/device/sriov_numvfs
```

3.5.1 Configure Open vSwitch hardware offloading

1. Unbind the VFs from the Mellanox driver:

```
# for vf in `grep PCI_SLOT_NAME /sys/class/net/ens4f0/device/virtfn*/
→uevent | cut -d '=' -f2`
do
  echo $vf > /sys/bus/pci/drivers/mlx5_core/unbind
done
```

2. Enable the switch in the NIC:

```
# PCI_ADDR=`grep PCI_SLOT_NAME /sys/class/net/ens4f0/device/uevent | sed
→'s:.*PCI_SLOT_NAME=:::'`
# devlink dev eswitch set pci/$PCI_ADDR mode switchdev
```

3. Enable hardware offload filters with TC:

```
# ethtool -K ens4f0 hw-tc-offload on
```

4. Rebind the VFs to the Mellanox driver:

```
# for vf in `grep PCI_SLOT_NAME /sys/class/net/ens4f0/device/virtfn*/
→uevent | cut -d '=' -f2`
do
    echo $vf > /sys/bus/pci/drivers/mlx5_core/bind
done
```

5. Enable hardware offloading in OVS:

```
# ovs-vsctl set Open_vSwitch . other_config:hw-offload=true
# ovs-vsctl set Open_vSwitch . other_config:max-idle=30000
```

6. Restart Open vSwitch

```
# systemctl restart openvswitch-switch
```

7. Restart the Open vSwitch agent

```
# systemctl restart neutron-openvswitch-agent
```

8. Restart the Nova compute service

```
# systemctl restart nova-compute
```

Warning: Changes to `sriov_numvfs` as well as the built-in NIC switch will not persist a reboot and must be performed every time the server is started.

3.6 Verify operation

To verify operation of hardware-offloaded Open vSwitch, you must create a virtual machine instance using an image with the proper network drivers.

The following images are known to contain working drivers:

- [Fedora 24](#)
- [Ubuntu 18.04 LTS \(Bionic\)](#)
- [Centos 7 \(1901\)](#)

Before creating an instance, a Neutron port must be created that has the following characteristics:

```
--vnic-type direct --binding-profile '{"capabilities":
["switchdev"]}'
```


To ensure flows are offloaded, disable port security with the `--disable-port-security` argument.

An example of the full command can be seen here:

```
# openstack port create \
  --network <network> \
  --vnic-type direct --binding-profile '{"capabilities": ["switchdev"]}' \
  --disable-port-security \
  <name>
```

The port can then be attached to the instance at boot. Once booted, the port will be updated to reflect the PCI address of the corresponding virtual function:

```
root@aio1-utility-container-8c0b0916:~# openstack port show -c binding_
↪profile testport2
+-----+-----+
↪ | Field          | Value                                     |
+-----+-----+
↪ | binding_profile | capabilities='[u'switchdev']', pci_slot='0000:21:00.6',
↪ pci_vendor_info='15b3:1016', physical_network='physnet1' |
+-----+-----+
```

3.6.1 Observing traffic

From the compute node, perform a packet capture on the representor port that corresponds to the virtual function attached to the instance. In this example, the interface is `eth1`.

```
root@computel1:~# tcpdump -nnn -i eth1 icmp
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth1, link-type EN10MB (Ethernet), capture size 262144 bytes
```

Perform a ping from another host and observe the traffic at the representor port:

```
root@infra2:~# ping 192.168.88.151 -c5
PING 192.168.88.151 (192.168.88.151) 56(84) bytes of data.
64 bytes from 192.168.88.151: icmp_seq=1 ttl=64 time=48.3 ms
64 bytes from 192.168.88.151: icmp_seq=2 ttl=64 time=1.52 ms
64 bytes from 192.168.88.151: icmp_seq=3 ttl=64 time=0.586 ms
64 bytes from 192.168.88.151: icmp_seq=4 ttl=64 time=0.688 ms
64 bytes from 192.168.88.151: icmp_seq=5 ttl=64 time=0.775 ms

--- 192.168.88.151 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4045ms
rtt min/avg/max/mdev = 0.586/10.381/48.335/18.979 ms

root@computel1:~# tcpdump -nnn -i eth1 icmp
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth1, link-type EN10MB (Ethernet), capture size 262144 bytes
19:51:09.684957 IP 192.168.88.254 > 192.168.88.151: ICMP echo request, id_
↪11168, seq 1, length 64
```

(continues on next page)

(continued from previous page)

```
19:51:09.685448 IP 192.168.88.151 > 192.168.88.254: ICMP echo reply, id ↵  
↵11168, seq 1, length 64
```

When offloading is handled in the NIC, only the first packet(s) of the flow will be visible in the packet capture.

The following command can be used to dump flows in the kernel datapath:

```
# ovs-dpctl dump-flows type=ovs
```

The following command can be used to dump flows that are offloaded:

```
# ovs-dpctl dump-flows type=offloaded
```

SCENARIO - USING OPEN VSWITCH WITH DVR

4.1 Overview

Operators can choose to utilize Open vSwitch with Distributed Virtual Routing (DVR) instead of Linux Bridges or plain Open vSwitch for the neutron ML2 agent. This offers the possibility to deploy virtual routing instances outside the usual neutron networking node. This document outlines how to set it up in your environment.

4.2 Recommended reading

This guide is a variation of the standard Open vSwitch deployment guide available at:

https://docs.openstack.org/openstack-ansible-os_neutron/latest/app-openvswitch.html

We recommend that you read the following documents before proceeding:

- Neutron documentation on Open vSwitch DVR OpenStack deployments: <https://docs.openstack.org/neutron/latest/admin/deploy-ovs-ha-dvr.html>
- Blog post on how OpenStack-Ansible works with Open vSwitch: <https://trumant.github.io/openstack-ansible-dvr-with-openvswitch.html>

4.3 Prerequisites

Configure your networking according the Open vSwitch setup:

- Scenario - Using Open vSwitch https://docs.openstack.org/openstack-ansible-os_neutron/latest/app-openvswitch.html

4.4 OpenStack-Ansible user variables

Create a group var file for your network hosts `/etc/openstack_deploy/group_vars/network_hosts`. It has to include:

```
# Ensure the openvswitch kernel module is loaded
openstack_host_specific_kernel_modules:
  - name: "openvswitch"
    pattern: "CONFIG_OPENVSWITCH"
```

Specify provider network definitions in your `/etc/openstack_deploy/openstack_user_config.yml` that define one or more Neutron provider bridges and related configuration:

Note: Bridges specified here will be created automatically. If `network_interface` is defined, the interface will be placed into the bridge automatically.

```
- network:
  container_bridge: "br-provider"
  container_type: "veth"
  type: "vlan"
  range: "101:200,301:400"
  net_name: "physnet1"
  network_interface: "bond1"
  group_binds:
    - neutron_openvswitch_agent
- network:
  container_bridge: "br-provider2"
  container_type: "veth"
  type: "vlan"
  range: "203:203,467:500"
  net_name: "physnet2"
  network_interface: "bond2"
  group_binds:
    - neutron_openvswitch_agent
```

When using `flat` provider networks, modify the network type accordingly:

```
- network:
  container_bridge: "br-provider"
  container_type: "veth"
  type: "flat"
  net_name: "flat"
  group_binds:
    - neutron_openvswitch_agent
```

Specify an overlay network definition in your `/etc/openstack_deploy/openstack_user_config.yml` that defines overlay network-related configuration:

Note: The bridge name should correspond to a pre-created Linux bridge or OVS bridge.

```
- network:
  container_bridge: "br-vxlan"
  container_type: "veth"
  container_interface: "eth10"
  ip_from_q: "tunnel"
  type: "vxlan"
  range: "1:1000"
  net_name: "vxlan"
  group_binds:
    - neutron_openvswitch_agent
```

Set the following user variables in your `/etc/openstack_deploy/user_variables.yml`:

Note: The only difference a DVR deployment and the standard Open vSwitch deployment is the setting of the respective `neutron_plugin_type`.

```
neutron_plugin_type: ml2.ovs.dvr
neutron_ml2_drivers_type: "flat,vlan,vxlan"
```

The overrides are instructing Ansible to deploy the OVS mechanism driver and associated OVS and DVR components. This is done by setting `neutron_plugin_type` to `ml2.ovs.dvr`.

The `neutron_ml2_drivers_type` override provides support for all common type drivers supported by OVS.

For additional information regarding provider network overrides and other configuration options, please refer to the standard Open vSwitch deployment available at:

https://docs.openstack.org/openstack-ansible-os_neutron/latest/app-openvswitch.html

SCENARIO - USING OPEN VSWITCH W/ DPDK

5.1 Overview

Operators can choose to utilize DPDK-accelerated Open vSwitch instead of unaccelerated Open vSwitch or Linux Bridges for the Neutron virtual network infrastructure. This architecture is best suited for NFV workloads and requires careful consideration and planning before implementing. This document outlines how to set it up in your environment.

Warning: The current implementation of DPDK in OpenStack-Ansible is experimental and not production ready. There is no guarantee of upgradability or backwards compatibility between releases.

5.2 Recommended reading

We recommend that you read the following documents before proceeding:

- Neutron with Open vSwitch Scenario: https://docs.openstack.org/openstack-ansible-os_neutron/latest/app-openvswitch.html
- Open vSwitch with DPDK datapath: <https://docs.openstack.org/neutron/latest/admin/config-ovs-dpdk.html>
- Getting the best performance from DPDK: https://doc.dpdk.org/guides-16.04/linux_gsg/nic_perf_intel_platform.html
- OpenStack documentation on hugepages: <https://docs.openstack.org/nova/latest/admin/huge-pages.html>

5.3 Prerequisites

To enable DPDK on a Linux platform, ensure that VT-d/VT-x are enabled for Intel processors and AMD-V/AMD-Vi are enabled for AMD processors. Such features are typically enabled in the BIOS.

On an Intel platform, the following kernel parameters are required and can be added to the GRUB configuration:

```
GRUB_CMDLINE_LINUX="... iommu=pt intel_iommu=on"
```

On an AMD platform, use these parameters instead:

```
GRUB_CMDLINE_LINUX="... iommu=pt amd_iommu=on"
```

Update GRUB and reboot the host(s).

Hugepages are required for DPDK. Instances leveraging DPDK-accelerated Open vSwitch must be configured to use hugepages by way of flavor attributes. Those attributes and the configuration of hugepages are described in this guide.

CPU frequency should be set to maximum for optimal performance. Many hardware vendors set the energy saving properties in the BIOS that may need to be modified. Changing the CPU frequency using `cpufreq` or similar utilities to performance from `ondemand` is recommended.

Note: The playbooks currently only support a single NIC interface for DPDK. Multiple ports per NIC are not yet supported but may be at a later time. This guide assumes the NIC is bound to NUMA node0, but the instructions can be modified for NICs bound to other NUMA nodes..

5.4 NUMA topology

Non-uniform memory access (NUMA) is a computer memory design used in multiprocessing. This guide cannot go into great depths about NUMA architecture. However, there are some configurations to be made that rely on the operator understanding NUMA characteristics of compute nodes hosting workloads using DPDK-accelerated Open vSwitch.

To view the NUMA topology of a particular system, use the `numactl` command shown here:

```
root@computel1:~# numactl --hardware
available: 2 nodes (0-1)
node 0 cpus: 0 1 2 3 4 5 6 7 16 17 18 19 20 21 22 23
node 0 size: 48329 MB
node 0 free: 31798 MB
node 1 cpus: 8 9 10 11 12 13 14 15 24 25 26 27 28 29 30 31
node 1 size: 48379 MB
node 1 free: 25995 MB
node distances:
node  0  1
  0:  10  20
  1:  20  10
```

The NUMA topology presented here corresponds to a host with 2x Intel Xeon 2450L processors with 96 GB of total RAM. The RAM is evenly split between the two NUMA nodes. Each CPU has 8 cores. With hyperthreading enabled, there are 16 threads per CPU for a total of 32 threads or cores presented to the operating system. It just so happens that this two-socket system has one NUMA node per socket, however, that will not always be the case. Consult your systems documentation for information unique to your system.

The first eight cores/cpus in the list for a given NUMA node can be considered physical cores in the CPU. For NUMA node0, this would be cores 0-7. The other eight cores, 16-23, are considered virtual sibling cores and are presented when hyperthreading is enabled. The physical-to-virtual mapping can be determined with the following commands:


```

root@compute1:~# for cpu in {0..7}; do cat /sys/devices/system/cpu/"cpu"
→$cpu/topology/thread_siblings_list; done
0,16
1,17
2,18
3,19
4,20
5,21
6,22
7,23

root@compute1:~# for cpu in {8..15}; do cat /sys/devices/system/cpu/"cpu"
→$cpu/topology/thread_siblings_list; done
8,24
9,25
10,26
11,27
12,28
13,29
14,30
15,31

```

A PCI slot typically corresponds to a single NUMA node. For optimal performance, a DPDK NIC and any instance utilizing the NIC should be restricted to the same NUMA node and its respective memory. Ensuring this behavior requires the use of flavors, host aggregates, and special kernel parameters and Open vSwitch/DPDK configuration settings.

In this example, a single 10G NIC installed in PCI slot 2 is bound to NUMA node0. Ideally, any instances utilizing the NIC would be limited to cores and memory associated with NUMA node0. This means cores 0-7 and 16-23, and up to 48GB of RAM. In reality, however, some cores and RAM from NUMA node0 will be reserved and made unavailable to instances. In addition, cores 8-15 and 24-31 associated with NUMA node1 should be made unavailable to instances. The configuration to do just that will be covered later in this guide.

It is considered good practice to reserve a single physical core and its respective virtual sibling from each NUMA node for normal (non-DPDK) operating system functions. In addition, at least one physical core (and sibling) from each NUMA node should be reserved for DPDK poll mode driver (PMD) functions, even when a NIC(s) is bound to a single NUMA node. The remaining cores can be reserved for virtual machine instances.

In this example, the breakdown would resemble the following:

Reserved Cores	Purpose	node0	node1
0,16,24	Host Operating System	0,16	8,24
1,9,17,25	DPDK PMDs	1,17	9,25
2-7,18-23	Virtual Machines	2-7,18-23	N/A

The variables and overrides used to define this configuration are discussed in the following sections.

5.5 Hugepage configuration

DPDK requires the configuration of hugepages, which is a mechanism by which the Linux kernel can partition and address larger amounts of memory beyond the basic page unit (4096 bytes). Huge pages are blocks of contiguous memory that commonly come in 2MB and 1G sizes. The page tables used by 2MB pages are suitable for managing multiple gigabytes of memory, whereas the page tables of 1GB pages are preferred for scaling to terabytes of memory. DPDK requires the use of 1GB pages.

A typical x86 system will have a Huge Page Size of 2048 kBytes (2MB). The default huge page size may be found by looking at the output of `/proc/meminfo`:

```
# cat /proc/meminfo | grep Hugepagesize
Hugepagesize: 2048 kB
```

The number of Hugepages can be allocated at runtime by modifying `/proc/sys/vm/nr_hugepages` or by using the `sysctl` command.

To view the current setting using the `/proc` entry:

```
# cat /proc/sys/vm/nr_hugepages
0
```

To view the current setting using the `sysctl` command:

```
# sysctl vm.nr_hugepages
vm.nr_hugepages = 0
```

To set the number of huge pages using `/proc` entry:

```
# echo 5 > /proc/sys/vm/nr_hugepages
```

To set the number of hugepages using `sysctl`:

```
# sysctl -w vm.nr_hugepages=5
vm.nr_hugepages = 5
```

It may be necessary to reboot to be able to allocate the number of hugepages that is needed. This is due to hugepages requiring large areas of contiguous physical memory.

When 1G hugepages are used, they must be configured at boot time. The amount of 1G hugepages that should be created will vary based on a few factors, including:

- The total amount of RAM available in the system
- The amount of RAM required for the planned number of instances
- The number of NUMA nodes that will be used

The NUMA topology presented here corresponds to a host with 2x Intel Xeon 2450L processors with 96GB of total RAM. The RAM is evenly split between the two NUMA nodes. A DPDK NIC will be associated with a single NUMA node, and for optimal performance any instance utilizing the DPDK NIC should be limited to the same cores and memory associated with the NUMA node. On this example system, both DPDK and instances can only utilize *up to* the 48GB of RAM associated with NUMA node0, though some of that RAM will be utilized by the OS and other tasks.

Of the 48GB of RAM available on NUMA node0, 32GB will be reserved for 1GB hugepages to be consumed by DPDK PMDs and instances. Configuring hugepages using kernel parameters results in

the defined number of hugepages to be split evenly across NUMA nodes. With the following kernel parameter, each NUMA node will be assigned 32x 1G hugepages:

```
GRUB_CMDLINE_LINUX="... hugepagesz=1G hugepages=64"
```

Hugepages can be adjusted at runtime if necessary, but doing so is outside the scope of this guide.

5.6 OpenStack-Ansible variables and overrides

The ability to pin instances to certain cores is not new, and can be accomplished using the `vcpu_pin_set` override seen here:

```
nova_nova_conf_overrides:
  DEFAULT:
    vcpu_pin_set: 2-7,18-23
```

This change can be added to the `user_overrides.yml` file for global implementation, or to individual nodes in the `openstack_user_config.yml` file as shown here:

```
compute_hosts:
  compute01:
    ip: 172.29.236.200
    container_vars:
      ...
      nova_nova_conf_overrides:
        DEFAULT:
          vcpu_pin_set: 2-7,18-23
```

Cores reserved for host operating system functions (non-DPDK) must be converted to a hexadecimal mask and defined using the `ovs_dpdk_lcore_mask` override. To convert to a hex mask you must first establish the binary mask of chosen cores using the following table:

```
31|30|. |24|23|. |17|16|15|. |9|8|7|. |1|0|
|1-|11-|111-|1111-|1111-|111
0|0|. |1|1|0|. |0|1|1|0|. |0|1|1|0|. |0|1|1|
```

The ellipses represent cores not shown. The binary mask for cores 0,8,16,24 can be determined in the following way:

```
00000001000000010000000100000001
```

The hexadecimal representation of that binary value is `0x1010101`. Set the `ovs_dpdk_lcore_mask` override accordingly in the `user_variables.yml` file or `openstack_user_config.yml`:

```
ovs_dpdk_lcore_mask: 1010101
```

The mask for cores 1,9,17,25 reserved for DPDK PMDs can be determined in a similar fashion. The table would resemble the following:

```
31|30|.|25|24|.|17|16|15|.|9|8|7|.|1|0|
| | - | | | - | | | | - | | | | - | | | |
0|0|.|1|1|0|.|1|1|0|0|.|1|1|0|0|.|1|1|0|
```

The ellipses represent cores not shown. The binary mask for cores 1,9,17,254 can be determined in the following way:

```
0000001000000001000000001000000010
```

The hexadecimal representation of that binary value is 0x2020202. Set the `ovs_dpdk_pmd_cpu_mask` override accordingly in the `user_variables.yml` file or `openstack_user_config.yml`:

```
ovs_dpdk_pmd_cpu_mask: 2020202
```

Additional variables should be set, including:

- `ovs_dpdk_driver`
- `ovs_dpdk_pci_addresses`
- `ovs_dpdk_socket_mem`

The default value for `ovs_dpdk_driver` is `vfio-pci`. Overrides can be set globally or on a per-host basis.

Note: Please consult the [DPDK Network Interface Controller Driver documentation](#) for more information on supported network drivers for DPDK.

The value for `ovs_dpdk_pci_addresses` is the PCI bus address of the NIC port(s) associated with the DPDK NIC. In this example, the DPDK NIC is identified as address `0000:03:00`. The individual interfaces are `0000:03:00.0` and `0000:03:00.1`, respectively. The variable `ovs_dpdk_pci_addresses` is a list, and both values can be defined like so:

```
ovs_dpdk_pci_addresses:
- 0000:03:00.0
- 0000:03:00.1
```

The value for `ovs_dpdk_socket_mem` will vary based on the number of NUMA nodes, number of NICs per NUMA node, and the MTU. The default value assumes a single NUMA node and associates a single 1G hugepage to DPDK that can handle a 1500 MTU. When multiple NUMA nodes are available, even with a single NIC, the following should be set:

```
ovs_dpdk_socket_mem: "1024,1024"
```

For systems using a single NUMA node of a dual-NUMA system and a 9000 MTU, the following can be set:

```
ovs_dpdk_socket_mem: "3072,1024"
```

Determining socket memory required involves calculations that are out of the scope of this guide.

5.7 Flavor configuration

Instances that connect to a DPDK-accelerated Open vSwitch must be configured to utilize large (1G) hugepages by way of custom flavor attributes.

The `hw:mem_page_size` property can be set on a new or existing flavor to enable this functionality:

```
openstack flavor set m1.small --property hw:mem_page_size=large
```

NOTE: If small page size is used, or no page size is set, the interface may appear in the instance but will not be functional.

5.8 OpenStack-Ansible user variables

Create a group var file for your network hosts `/etc/openstack_deploy/group_vars/network_hosts`. It has to include:

```
# Ensure the openvswitch kernel module is loaded
openstack_host_specific_kernel_modules:
  - name: "openvswitch"
    pattern: "CONFIG_OPENVSWITCH"
```

Specify provider network definitions in your `/etc/openstack_deploy/openstack_user_config.yml` that define one or more Neutron provider bridges and related configuration:

```
- network:
  container_bridge: "br-provider"
  container_type: "veth"
  type: "vlan"
  range: "101:200,301:400"
  net_name: "physnet1"
  group_binds:
    - neutron_openvswitch_agent
```

Note: A single DPDK interface can be connected to an OVS provider bridge, and must be done using the `ovs-vsctl` command as a post-installation step.

Set the following user variables in your `/etc/openstack_deploy/user_variables.yml` to enable the Open vSwitch driver and DPDK support:

```
neutron_plugin_type: ml2.ovs
neutron_ml2_drivers_type: "vlan"

# Enable DPDK support
ovs_dpdk_support: True

# Add these overrides or set on per-host basis in openstack_user_config.yml
ovs_dpdk_pci_addresses: "0000:03:00.0"
ovs_dpdk_lcore_mask: 1010101
ovs_dpdk_pmd_cpu_mask: 2020202
ovs_dpdk_socket_mem: "1024,1024"
```

Note: Overlay networks are not supported on DPDK-enabled nodes at this time.

5.9 Post-installation

Once the playbooks have been run and OVS/DPDK has been configured, it will be necessary to add a physical interface to the provider bridge before networking can be fully established.

On compute nodes, the following command can be used to attach a NIC port 0000:03:00.0 to the provider bridge br-provider:

```
ovs-vsctl add-port br-provider 0000:03:00.0 -- set Interface 0000:03:00.0
↳type=dtpk options:dtpk-devargs=0000:03:00.0
```

The command can be adjusted according to your configuration.

Warning: Adding multiple ports to the bridge may result in bridging loops unless bonding is configured. DPDK bonding is outside the scope of this guide.

SCENARIO - USING OPEN VSWITCH WITH SFC

6.1 Overview

Operators can choose to configure SFC mechanism with Open vSwitch instead of ODL through the neutron networking-sfc project. The SFC configuration results in OVS flows being configured with SFC specifics using MPLS as dataplane technology. This document outlines how to set it up in your environment.

6.2 Recommended reading

We recommend that you read the following documents before proceeding:

- General overview of neutron networking-sfc: <https://docs.openstack.org/networking-sfc/latest/index.html>
- How to configure networking-sfc in neutron: <https://docs.openstack.org/networking-sfc/latest/install/configuration.html>

6.3 Prerequisites

Configure your networking according the Open vSwitch setup:

- Scenario - Using Open vSwitch https://docs.openstack.org/openstack-ansible-os_neutron/latest/app-openvswitch.html

6.4 OpenStack-Ansible user variables

Create a group var file for your network hosts `/etc/openstack_deploy/group_vars/network_hosts`. It has to include:

```
# Ensure the openvswitch kernel module is loaded
openstack_host_specific_kernel_modules:
  - name: "openvswitch"
    pattern: "CONFIG_OPENVSWITCH"
```

Set the following user variables in your `/etc/openstack_deploy/user_variables.yml`:

```
### neutron specific config
neutron_plugin_type: ml2.ovs

neutron_ml2_drivers_type: "flat,vlan"

neutron_plugin_base:
- router
- metering
- flow_classifier
- sfc

# Typically this would be defined by the os-neutron-install
# playbook. The provider_networks library would parse the
# provider_networks list in openstack_user_config.yml and
# generate the values of network_types, network_vlan_ranges
# and network_mappings. network_mappings would have a
# different value for each host in the inventory based on
# whether or not the host was metal (typically a compute host)
# or a container (typically a neutron agent container)
#
# When using Open vSwitch, we override it to take into account
# the Open vSwitch bridge we are going to define outside of
# OpenStack-Ansible plays
neutron_provider_networks:
  network_flat_networks: "*"
  network_types: "vlan"
  network_vlan_ranges: "physnet1:102:199"
  network_mappings: "physnet1:br-provider"
```

Note: The only difference to the Standard Open vSwitch configuration is the setting of the neutron_plugin_base.

SCENARIO - OPEN VIRTUAL NETWORK (OVN)

7.1 Overview

Operators can choose to utilize the Open Virtual Network (OVN) mechanism driver instead of Linux bridges or plain Open vSwitch for the Neutron ML2 plugin. This offers the possibility to deploy virtual networks and routers using OVN with Open vSwitch, which replaces the agent-based model used by the aforementioned architectures. This document outlines how to set it up in your environment.

Warning: The current implementation of OVN in OpenStack-Ansible is experimental and not production ready.

The following architectural assumptions have been made:

- Each compute node will act as an OVN controller
- Each compute node is eligible to serve as an OVN gateway node

Note: Physical VTEP integration is not supported.

7.2 Recommended reading

Since this is an extension of the basic Open vSwitch scenario, it is worth reading that scenario to get some background. It is also recommended to be familiar with OVN and networking-ovn projects and their configuration.

- [Scenario: Open vSwitch](#)
- [OVN Architecture](#)
- [Networking-ovn](#)

7.3 Prerequisites

- Open vSwitch >= 2.10.0
- A successful deployment of OVN requires a dedicated network interface be attached to the OVS provider bridge. This is not handled automatically and may require changes to the network interface configuration file.

7.4 OpenStack-Ansible user variables

Create a group var file for your network hosts `/etc/openstack_deploy/group_vars/network_hosts`. It has to include:

```
# Ensure the openvswitch kernel module is loaded
openstack_host_specific_kernel_modules:
  - name: "openvswitch"
    pattern: "CONFIG_OPENVSWITCH"
```

Copy the neutron environment overrides to `/etc/openstack_deploy/env.d/neutron.yml` to disable the creation of the neutron agents container and implement the `neutron_ovn_northd_container` hosts group containing all network nodes:

```
component_skel:
  neutron_ovn_controller:
    belongs_to:
      - neutron_all
  neutron_ovn_northd:
    belongs_to:
      - neutron_all

container_skel:
  neutron_agents_container:
    contains: {}
  neutron_ovn_northd_container:
    belongs_to:
      - network_containers
    contains:
      - neutron_ovn_northd
```

Copy the nova environment overrides to `/etc/openstack_deploy/env.d/nova.yml` to implement the `neutron_ovn_controller` hosts group containing all compute nodes:

```
container_skel:
  nova_compute_container:
    belongs_to:
      - compute_containers
      - kvm-compute_containers
      - lxd-compute_containers
      - qemu-compute_containers
    contains:
      - neutron_ovn_controller
      - nova_compute
  properties:
    is_metal: true
```

Specify provider network definitions in your `/etc/openstack_deploy/openstack_user_config.yml` that define one or more Neutron provider bridges and related configuration:

Note: Bridges specified here will be created automatically. If `network_interface` is defined, the interface will be placed into the bridge automatically. Only VLAN network types are supported at this time.

```
- network:
  container_bridge: "br-privatenet"
  container_type: "veth"
  type: "vlan"
  range: "101:200,301:400"
  net_name: "private"
  network_interface: "bond2"
  group_binds:
    - neutron_ovn_controller
- network:
  container_bridge: "br-publicnet"
  container_type: "veth"
  type: "vlan"
  range: "203:203,467:500"
  net_name: "public"
  network_interface: "bond1"
  group_binds:
    - neutron_ovn_controller
```

Specify an overlay network definition in your `/etc/openstack_deploy/openstack_user_config.yml` that defines overlay network-related configuration:

Note: The bridge name should correspond to a pre-created Linux bridge. Only GENEVE overlay network types are supported at this time.

```
- network:
  container_bridge: "br-vxlan"
  container_type: "veth"
  container_interface: "eth10"
  ip_from_q: "tunnel"
  type: "geneve"
  range: "1:1000"
  net_name: "geneve"
  group_binds:
    - neutron_ovn_controller
```

Set the following user variables in your `/etc/openstack_deploy/user_variables.yml`:

```
neutron_plugin_type: ml2.ovn

neutron_plugin_base:
  - neutron.services.ovn_l3.plugin.OVNL3RouterPlugin

neutron_ml2_drivers_type: "vlan,local,geneve"
```

The overrides are instructing Ansible to deploy the OVN mechanism driver and associated OVN components. This is done by setting `neutron_plugin_type` to `ml2.ovn`.

The `neutron_plugin_base` override instructs Neutron to use OVN for routing functions rather than the standard L3 agent model.

The `neutron_ml2_drivers_type` override provides support for all type drivers supported by OVN.

If provider network overrides are needed on a global or per-host basis, the following format can be used in `user_variables.yml` or per-host in `openstack_user_config.yml`.

Note: These overrides are not normally required.

```
# When configuring Neutron to support geneve tenant networks and
# vlan provider networks the configuration may resemble the following:
neutron_provider_networks:
  network_types: "geneve"
  network_geneve_ranges: "1:1000"
  network_vlan_ranges: "public"
  network_mappings: "public:br-publicnet"
  network_interface_mappings: "br-publicnet:bond1"

# When configuring Neutron to support only vlan tenant networks and
# vlan provider networks the configuration may resemble the following:
neutron_provider_networks:
  network_types: "vlan"
  network_vlan_ranges: "public:203:203,467:500"
  network_mappings: "public:br-publicnet"
  network_interface_mappings: "br-publicnet:bond1"

# When configuring Neutron to support multiple vlan provider networks
# the configuration may resemble the following:
neutron_provider_networks:
  network_types: "vlan"
  network_vlan_ranges: "public:203:203,467:500,private:101:200,301:400"
  network_mappings: "public:br-publicnet,private:br-privatenet"
  network_interface_mappings: "br-publicnet:bond1,br-privatenet:bond2"
```

7.5 (Optional) DVR or Distributed L3 routing

DVR will be used for floating IPs if the `ovn / enable_distributed_floating_ip` flag is configured to `True` in the neutron server configuration.

Create a group var file for neutron server `/etc/openstack_deploy/group_vars/neutron_server.yml`. It has to include:

```
# DVR/Distributed L3 routing support
neutron_neutron_conf_overrides:
  ovn:
    enable_distributed_floating_ip: True
```

7.6 Open Virtual Network (OVN) commands

The following commands can be used to provide useful information about the state of Open vSwitch networking and configurations.

The `ovs-vsctl list open_vswitch` command provides information about the `open_vswitch` table in the local Open vSwitch database:

```
root@aio1:~# ovs-vsctl list open_vswitch
_uuid          : 855c820b-c082-4d8f-9828-8cab01c6c9a0
bridges        : [37d3bd82-d436-474e-89b7-705aea634d7d, a393b2f6-5c3d-
↳4ccd-a2f9-e9817391612a]
cur_cfg        : 14
datapath_types : [netdev, system]
db_version     : "7.15.1"
external_ids   : {hostname="aio1", ovn-bridge-mappings="vlan:br-
↳provider", ovn-encap-ip="172.29.240.100", ovn-encap-type="geneve,vxlan",
↳ovn-remote="tcp:172.29.236.100:6642", rundir="/var/run/openvswitch",
↳system-id="11af26c6-9ec1-4cf7-bf41-2af45bd59b03"}
iface_types    : [geneve, gre, internal, lisp, patch, stt, system,
↳tap, vxlan]
manager_options : []
next_cfg       : 14
other_config   : {}
ovs_version    : "2.9.0"
ssl            : []
statistics     : {}
system_type    : ubuntu
system_version : "16.04"
```

The `ovn-sbctl show` command provides information related to southbound connections. If used outside the `ovn_northd` container, specify the connection details:

```
root@aio1-neutron-ovn-northd-container-57a6f1a9:~# ovn-sbctl show
Chassis "11af26c6-9ec1-4cf7-bf41-2af45bd59b03"
  hostname: "aio1"
  Encap vxlan
    ip: "172.29.240.100"
    options: {csum="true"}
  Encap geneve
    ip: "172.29.240.100"
    options: {csum="true"}

root@aio1:~# ovn-sbctl --db=tcp:172.29.236.100:6642 show
Chassis "11af26c6-9ec1-4cf7-bf41-2af45bd59b03"
  hostname: "aio1"
  Encap vxlan
    ip: "172.29.240.100"
    options: {csum="true"}
  Encap geneve
    ip: "172.29.240.100"
    options: {csum="true"}
```

The `ovn-nbctl show` command provides information about networks known to OVN and demonstrates connectivity between the northbound database and neutron-server.

```
root@aio1-neutron-ovn-northd-container-57a6f1a9:~# ovn-nbctl show
switch 5e77f29e-5dd3-4875-984f-94bd30a12dc3 (neutron-87ec5a05-9abe-4c93-
↪89bd-c6d40320db87) (aka testnet)
  port 65785045-69ec-49e7-82e3-b9989f718a9c
    type: localport
    addresses: ["fa:16:3e:68:a3:c8"]
```

The `ovn-nbctl list Address_Set` command provides information related to security groups. If used outside the `ovn_northd` container, specify the connection details:

```
root@aio1-neutron-ovn-northd-container-57a6f1a9:~# ovn-nbctl list Address_
↪Set
_uuid          : 575b3015-f83f-4bd6-a698-3fe67e43bec6
addresses      : []
external_ids   : {"neutron:security_group_id"="199997c1-6f06-4765-
↪89af-6fd064365c6a"}
name           : "as_ip4_199997c1_6f06_4765_89af_6fd064365c6a"

_uuid          : b6e211af-e52e-4c59-93ce-adf75ec14f46
addresses      : []
external_ids   : {"neutron:security_group_id"="199997c1-6f06-4765-
↪89af-6fd064365c6a"}
name           : "as_ip6_199997c1_6f06_4765_89af_6fd064365c6a"

root@aio1:~# ovn-nbctl --db=tcp:172.29.236.100:6641 list Address_Set
_uuid          : 575b3015-f83f-4bd6-a698-3fe67e43bec6
addresses      : []
external_ids   : {"neutron:security_group_id"="199997c1-6f06-4765-
↪89af-6fd064365c6a"}
name           : "as_ip4_199997c1_6f06_4765_89af_6fd064365c6a"

_uuid          : b6e211af-e52e-4c59-93ce-adf75ec14f46
addresses      : []
external_ids   : {"neutron:security_group_id"="199997c1-6f06-4765-
↪89af-6fd064365c6a"}
name           : "as_ip6_199997c1_6f06_4765_89af_6fd064365c6a"
```

Additional commands can be found in upstream OVN documentation.

7.7 Notes

The `ovn-controller` service on compute nodes will check in as an agent and can be observed using the `openstack network agent list` command:

```
root@aio1-utility-container-35bebd2a:~# openstack network agent list
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
↪+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | Agent Type |
↪Host | Availability Zone | Alive | State | Binary |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
↪+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 4db288a6-8f8a-4153-b4b7-7eaf44f9e881 | OVN Controller Gateway agent |
↪aio1 | n/a | :- ) | UP | ovn-controller |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
↪+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

(continues on next page)

(continued from previous page)

The HAproxy implementation in use may not properly handle active/backup failover for ovsdb-server with OVN. Work may be done to implement pacemaker/corosync or wait for upstream active/active support.

SCENARIO - USING THE NUAGE NEUTRON PLUGIN

8.1 Introduction

Follow these steps to deploy Nuage Networks VCS with OpenStack-Ansible (OSA):

- Install prerequisites.
- Configure neutron to use the Nuage Networks neutron plugin.
- Configure the Nuage Networks neutron plugin.
- Download Nuage Networks VCS components and playbooks.
- Execute the playbooks.

8.2 Prerequisites

1. The deployment environment is configured according to OSA best practices such as cloning OSA software and bootstrapping Ansible. See [OpenStack-Ansible Install Guide](#).
2. VCS stand-alone components, VSD and VSC, are configured and deployed. See the Nuage Networks VSD and VSC Install Guides.
3. Nuage VRS playbooks were cloned to the deployment host from <https://github.com/nuagenetworks/nuage-openstack-ansible>. This guide assumes a deployment host path of `/opt/nuage-openstack-ansible`.

8.3 Configure Nuage neutron plugin

Configuring the neutron plugin requires creating or editing of parameters in the following two files:

- `/etc/openstack_deploy/user_nuage_vars.yml`
- `/etc/openstack_deploy/user_variables.yml`

On the deployment host, copy the Nuage user variables file from `/opt/nuage-openstack-ansible/etc/user_nuage_vars.yml` to the `/etc/openstack_deploy/` folder.

```
# cp /opt/nuage-openstack-ansible/etc/user_nuage_vars.yml \
/etc/openstack_deploy/
```

Next, modify the following parameters in that file as per your Nuage VCS environment:

1. Replace *VSD Enterprise Name* with the name of VSD Enterprise:

```
nuage_net_partition_name: "<VSD Enterprise Name>"
```

2. Replace *VSD IP* and *VSD GUI Port* as per your VSD configuration:

```
nuage_vsd_ip: "<VSD IP>:<VSD GUI Port>"
```

3. Replace *VSD Username*, *VSD Password*, and *VSD Organization Name* with your login credentials for the VSD GUI:

```
nuage_vsd_username: "<VSD Username>"
nuage_vsd_password: "<VSD Password>"
nuage_vsd_organization: "<VSD Organization Name>"
```

4. Replace *Nuage VSP Version* with the Nuage VSP release for Integration. For example, for Nuage VSP release 3.2 this value would be *v3_2*.

```
nuage_base_uri_version: "<Nuage VSP Version>"
```

5. Replace *Nuage VSD CMS Id* with the CMS-Id generated by VSD to manage your OpenStack cluster:

```
nuage_cms_id: "<Nuage VSD CMS Id>"
```

6. Replace *Active VSC-IP* with the IP address of your active VSC node and *Standby VSC-IP* with the IP address of your standby VSC node:

```
active_controller: "<Active VSC-IP>"
standby_controller: "<Standby VSC-IP>"
```

7. Replace *Local Package Repository* with the link of the local repository hosting the Nuage VRS packages. For example: `http://192.0.2.10/debs/3.2/vrs/`:

```
nuage_vrs_debs_repo: "deb <Local Package Repository>"
```

8. On the Deployment host, add the following lines to the `/etc/openstack_deploy/user_variables.yml` file, replacing the *Local PyPi Mirror URL* with the link to the PyPi server hosting the Nuage OpenStack Python packages in `.whl` format:

```
neutron_plugin_type: "nuage"
nova_network_type: "nuage"
pip_links:
  - { name: "openstack_release", link: "{{ openstack_repo_url \
    }}/os-releases/{{ openstack_release }}/"}
  - { name: "nuage_repo", link: "<Local PyPi Mirror URL>" }
```

8.4 Installation

1. After you set up the multi-node OpenStack cluster, start the OpenStack deployment as listed in the OpenStack-Ansible Install guide by running all playbooks in sequence on the deployment host.
2. After OpenStack deployment is complete, deploy Nuage VRS on all compute target hosts in the OpenStack cluster by running the Nuage VRS playbooks in `/opt/nuage-openstack-ansible/nuage_playbook` on your deployment host:

```
# cd /opt/nuage-openstack-ansible/nuage_playbooks
# openstack-ansible nuage_all.yml
```

Note: To obtain Nuage Networks VSP software packages, user documentation, and licenses, contact `info@nuagenetworks.net`.

SCENARIO - VMWARE NSX PLUGIN

9.1 Introduction

This document covers the steps to integrate the VMware NSX plugin with OpenStack Ansible.

Warning: Currently, only NSX-T Policy API is supported.

Please follow these steps:

- Configure Neutron to use the NSX plugin

9.2 Prerequisites

1. The deployment environment is configured according to OSA best practices such as cloning OSA software and bootstrapping Ansible. See [OpenStack-Ansible Install Guide](#).
2. NSX-T has been deployed per its installation guide and compute nodes have been properly configured as transport nodes. See *NSX-T Data Center Installation Guide* <<https://docs.vmware.com/en/VMware-NSX-T-Data-Center/3.0/installation/GUID-3E0C4CEC-D593-4395-84C4-150CD6285963.htm>> _.

9.3 Configure Neutron to use the NSX plugin

Copy the neutron environment overrides to `/etc/openstack_deploy/env.d/neutron.yml` and disable agent creation, since it is not needed.

```
neutron_agents_container:  
  belongs_to:  
    - network_containers  
  contains: { }
```

Copy the nova environment overrides to `/etc/openstack_deploy/env.d/nova.yml` and disable neutron agent creation, since it is not needed.

```
container_skel:  
  nova_api_container:  
    belongs_to:
```

(continues on next page)

(continued from previous page)

```

- compute-infra_containers
- os-infra_containers
contains:
- nova_api_metadata
- nova_api_os_compute
- nova_conductor
- nova_scheduler
- nova_console
nova_compute_container:
  belongs_to:
  - compute_containers
  - kvm-compute_containers
  - qemu-compute_containers
contains:
- nova_compute
properties:
  is_metal: true

```

Set the following required variables in your `/etc/openstack_deploy/user_variables.yml`

```

neutron_plugin_type: vmware.nsx
nova_network_type: nsx
nsx_api_password: <password>
nsx_api_managers:
- nsx-manager-01
- nsx-manager-02
- nsx-manager-03

```

Optionally specify additional parameters using overrides

```

neutron_nsx_conf_ini_overrides:
  nsx_p:
  default_tier0_router: my-tier0-router
  default_overlay_tz: my-overlay-tz
  default_vlan_tz: my-vlan-tz
  metadata_proxy: my-metadata-proxy-profile
  dhcp_profile: my-dhcp-profile

```

Warning: If NSX has defined more than one tier 0, overlay/vlan tz, metadata proxy, or dhcp profile, then you must explicitly define those using conf overrides. Neutron will fail to start if these are not defined in those conditions.

9.4 Installation

After the environment has been configured as detailed above, start the OpenStack deployment as listed in the OpenStack-Ansible Install Guide.

SCENARIO - USING THE NETWORKING-CALICO NEUTRON PLUGIN

10.1 Introduction

This document describes the steps required to deploy Project Calico Neutron networking with OpenStack-Ansible (OSA). These steps include:

- Configure OSA environment overrides.
- Configure OSA user variables.
- Execute the playbooks.

For additional configuration about Project Calico and its architecture, please reference the [networking-calico](#) and [Project Calico](#) documentation.

10.2 Prerequisites

1. The deployment environment has been configured according to OSA best-practices. This includes cloning OSA software and bootstrapping Ansible. See [OpenStack-Ansible Install Guide](#)
2. BGP peers configured to accept routing announcements from your hypervisors. By default, the hypervisors default router is set as the BGP peer.

10.3 Configure OSA Environment for Project Calico

Add hosts to the `/etc/openstack_deploy/conf.d/etcd.conf` configuration file to add container hosts for the etcd cluster. See `etc/openstack_deploy/conf.d/etcd.conf.example` in the `openstack-ansible` repo or adjust the example below to match your infrastructure hosts:

```
etcd_hosts:
  infra1:
    ip: 172.20.236.111
  infra2:
    ip: 172.20.236.112
  infra3:
    ip: 172.20.236.113
```

Copy the neutron environment overrides to `/etc/openstack_deploy/env.d/neutron.yml` to disable the creation of the neutron agents container, and implement the calico-dhcp-agent hosts group containing all compute hosts.

```
component_skel:
  neutron_calico_dhcp_agent:
    belongs_to:
      - neutron_all

container_skel:
  neutron_agents_container:
    contains: {}
  neutron_calico_dhcp_agent_container:
    belongs_to:
      - compute_containers
    contains:
      - neutron_calico_dhcp_agent
  properties:
    is_metal: true
    service_name: neutron
```

10.4 Configure networking-calico Neutron Plugin

Set the following in `/etc/openstack_deploy/user_variables.yml`.

```
neutron_plugin_type: ml2.calico
nova_network_type: calico
```

10.5 Installation

After multi-node OpenStack cluster is configured as detailed above; start the OpenStack deployment as listed in the OpenStack-Ansible Install guide by running all playbooks in sequence on the deployment host

SCENARIO - OPENDAYLIGHT AND OPEN VSWITCH

11.1 Overview

Deployers can choose to enhance neutron capabilities by means of the OpenDaylight SDN Controller, which works together with Open vSwitch to provide advanced networking capabilities. This document explains how to use them in your environment.

11.2 Recommended reading

Since this is an extension of the basic Open vSwitch scenario, it is worth reading that scenario to get some background. It is also recommended to be familiar with OpenDaylight and networking-odl projects and their configuration.

- Scenario: Open vSwitch
- OpenDaylight SDN Controller
- Networking-odl

11.3 Prerequisites

The OpenDaylight Ansible role needs to be available in Ansibles role path.

11.4 OpenStack-Ansible user variables

Set the following user variables in your `/etc/openstack_deploy/user_variables.yml`:

```
### Ensure the openvswitch kernel module is loaded
openstack_host_specific_kernel_modules:
  - name: "openvswitch"
    pattern: "CONFIG_OPENVSWITCH"
    group: "network_hosts"

### Use OpenDaylight SDN Controller
neutron_plugin_type: "ml2.opendaylight"
odl_ip: "{{ hostvars[groups['opendaylight'][0]]['ansible_facts']['default_
↳ipv4']]['address'] }}"
neutron_opendaylight_conf_ini_overrides:
```

(continues on next page)

(continued from previous page)

```
m12_odl:
  url: "http://{{ odl_ip }}:8180/controller/nb/v2/neutron"
  username: <username>
  password: <password>
```

Most of the content of this file is self-explanatory. The first block is used to deploy Open vSwitch in all network hosts.

The second block is instructing Ansible to deploy OpenDaylight SDN Controller. This is done by specifying `neutron_plugin_type` to `m12.opendaylight`. The IP address of the OpenDaylight controller needs to be inferred from the deployment configuration as well. That can be used with a line such as the one in the example.

After that, some configuration is needed to integrate OpenDaylight and Neutron, using the `m12_odl` section.

- **url:** OpenDaylight's northbound url. This is automatically retrieved from the deployment configuration, so just need to copy the example line.
- **username:** OpenDaylight northbound API username
- **password:** OpenDaylight northbound API password for `<username>`

Apart from these options, the deployer might want to change the installation method for OpenDaylight Ansible role. This role uses pre-packaged binaries, which can be either `deb` or `rpm` files, and by default it will download these binaries from OpenDaylight repositories, trying to guess the correct package depending on the underlying operating system.

Also, the set of features that will be enabled in the OpenDaylight SDN controller defaults to `odl-netvirt-openstack`, which is the minimum for an OpenStack integration. The deployer can modify this value by providing a list of feature names in the `.opendaylight_extra_features` variable.

For more information, see OpenDaylight Ansible role documentation.

11.5 L3 configuration

L3 services are by default provided by the `neutron-l3-agent`. ODL is capable of providing L3 services too and if ODL is deployed, it is actually recommended to use them instead of `neutron`. Remember that L3 services allow, among other things, to give VMs connectivity to the internet.

To activate the ODL L3 services, you should add to the above explained variables:

```
# Activate the L3 capabilities of ODL
neutron_plugin_base:
  - odl-router_v2
  - metering
```

If you want to use the L3 capabilities, you will need to define an external Neutron network and set a gateway. Note that the `br-vlan` interface of the nodes could be a perfect interface for that gateway, although it depends on your network topology.

11.6 SFC configuration

It is possible to have an openstack-ansible deployment with SFC capabilities. The following config needs to be added to the above described `/etc/openstack_deploy/user_variables.yml`:

```
neutron_plugin_base:  
- router  
- metering  
- flow_classifier  
- sfc
```

When using this configuration, `networking-sfc` will be deployed and SFC features will be activated in ODL. A SFC topology could be then set up through the `networking-sfc` API or through an orchestrator like `tacker` (if deployed).

11.7 BGPVPN configuration

ODL provides support for extending L3 services over DC-GW by BGPVPN. This way Openstack configures ODL as BGP speaker to exchange the routes with DC-GW to establish the communication between Tenant VMs and external world in the data path.

To activate BGPVPN service, you should add the following variables in addition to the OpenStack-Ansible user variables mentioned above.

```
# Activate the BGPVPN capabilities of ODL  
neutron_plugin_base:  
- odl-router_v2  
- bgpvpn
```

11.8 Security information

Communications between the OpenDaylight SDN Controller and Open vSwitch are not secured by default. For further information on securing this interface, see these manuals:

- [TLS Support on OpenDaylight OpenFlow plugin](#).
- [Secure Communication Between OpenFlow Switches and Controllers](#).

SCENARIO - NETWORKING GENERIC SWITCH

12.1 Overview

Operators can choose to utilize the Networking Generic Switch (NGS) mechanism driver to manage physical switches when Ironic is integrated with Neutron. The Networking Generic Switch mechanism driver can be deployed alongside other drivers, such as Open vSwitch or LinuxBridge. This document outlines how to set it up in your environment.

12.2 Recommended reading

It is recommended to familiarize yourself with project-specific documentation to better understand deployment and configuration options:

- [Networking Generic Switch](#)

12.3 Prerequisites

- [Ironic Bare-Metal Provisioning Service](#)
- [Supported Network Hardware](#)
- Network connectivity from the node(s) running the *neutron-server* service to the management interface of the physical switch(es) connected to Ironic bare-metal nodes. This is outside the scope of OpenStack-Ansible.

12.4 OpenStack-Ansible user variables

Add `ml2.genericswitch` to the `neutron_plugin_types` list in `/etc/openstack_deploy/user_variables.yml`:

```
neutron_plugin_types:  
- ml2.genericswitch
```

To interface with a supported network switch, configure ini overrides for each connected switch in your environment:

neutron_ml2_conf_genericswitch_ini_overrides:

```
genericswitch:arista01:
  device_type: netmiko_arista_eos
  ngs_mac_address: "00:1c:73:29:ea:ca"
  ip: "192.168.90.2"
  username: "openstack"
  password: "0p3nst@ck"
  ngs_port_default_vlan: 3
genericswitch:arista02:
  device_type: netmiko_arista_eos
  ngs_mac_address: "00:1c:73:29:ea:cb"
  ip: "192.168.90.3"
  username: "openstack"
  password: "0p3nst@ck"
  ngs_port_default_vlan: 3
```

Lastly, configure an override to Ironic to enable the neutron interface:

```
ironic_enabled_network_interfaces_list: neutron
ironic_default_network_interface: neutron
```

12.5 Notes

Ironic bare-metal ports that are associated with bare-metal nodes can be configured with the respective connection details using the `openstack baremetal port set` command:

```
openstack baremetal port set 3a948c3b-6c41-4f68-8389-c4f5ca667c63 \
--local-link-connection switch_info=arista01 \
--local-link-connection switch_id="00:1c:73:29:ea:ca" \
--local-link-connection port_id="et11"
```

When a server is deployed using a bare-metal node, Neutron will connect to the respective switch(es) and configure the switchport interface(s) according.

tags openstack, neutron, cloud, ansible

category *nix

This role installs the following Systemd services:

- neutron-server
- neutron-agents

To clone or view the source code for this repository, visit the role repository for [os_neutron](#).

DEFAULT VARIABLES

```
###  
### Verbose Options  
###  
  
debug: False  
  
###  
### Service setup options  
###  
  
# Set the host which will execute the shade modules  
# for the service setup. The host must already have  
# clouds.yaml properly configured.  
neutron_service_setup_host: "{{ openstack_service_setup_host | default(  
→'localhost') }}"  
neutron_service_setup_host_python_interpreter: "{{ openstack_service_setup_  
→host_python_interpreter | default((neutron_service_setup_host ==  
→'localhost') | ternary(ansible_playbook_python, ansible_facts['python'] [  
→'executable'])) }}"  
  
###  
### Packages Options  
###  
  
# Set the package install state for distribution  
# Options are 'present' and 'latest'  
neutron_package_state: "{{ package_state | default('latest') }}"  
  
# Set installation method.  
neutron_install_method: "{{ service_install_method | default('source') }}"  
neutron_venv_python_executable: "{{ openstack_venv_python_executable |  
→default('python3') }}"  
  
###  
### Python code details  
###  
  
# Set the package install state for pip_package  
# Options are 'present' and 'latest'  
neutron_pip_package_state: "latest"  
  
# Source git repo/branch settings
```

(continues on next page)

(continued from previous page)

```

neutron_git_repo: https://opendev.org/openstack/neutron
neutron_git_install_branch: master
neutron_vpnaas_git_repo: https://opendev.org/openstack/neutron-vpnaas
neutron_vpnaas_git_install_branch: master
neutron_dynamic_routing_git_repo: https://opendev.org/openstack/neutron-
↳dynamic-routing
neutron_dynamic_routing_git_install_branch: master
networking_calico_git_repo: https://github.com/projectcalico/networking-
↳calico
networking_calico_git_install_branch: master
networking_odl_git_repo: https://opendev.org/openstack/networking-odl
networking_odl_git_install_branch: master
networking_sfc_git_repo: https://opendev.org/openstack/networking-sfc
networking_sfc_git_install_branch: master
networking_bgpvpn_git_repo: https://opendev.org/openstack/networking-bgpvpn
networking_bgpvpn_git_install_branch: master
ceilometer_git_repo: https://opendev.org/openstack/ceilometer
ceilometer_git_install_branch: master
networking_generic_switch_git_repo: https://opendev.org/openstack/
↳networking-generic-switch
networking_generic_switch_git_install_branch: master
networking_nsx_git_repo: https://opendev.org/x/vmware-nsx
networking_nsx_git_install_branch: master
networking_nsxlib_git_repo: https://opendev.org/x/vmware-nsxlib
networking_nsxlib_git_install_branch: master

neutron_upper_constraints_url: "{{ requirements_git_url | default('https://
↳releases.openstack.org/constraints/upper/' ~ requirements_git_install_
↳branch | default('master')) }}"
neutron_git_constraints:
  - "--constraint {{ neutron_upper_constraints_url }}"

neutron_pip_install_args: "{{ pip_install_options | default('') }}"

# Name of the virtual env to deploy into
neutron_venv_tag: "{{ venv_tag | default('untagged') }}"

###
### Generic Neutron Config
###

# Fatal Deprecations
neutron_fatal_deprecations: False

# If ``neutron_api_workers`` is unset the system will use half the number_
↳of available VCPUs to
# compute the number of api workers to use with a default capping value of_
↳16.
# neutron_api_workers: 16

## Cap the maximum number of threads / workers when a user value is_
↳unspecified.
neutron_api_threads_max: 16
neutron_api_threads: "{{ [[ansible_facts['processor_vcpus']]|default(2) //_
↳2, 1] | max, neutron_api_threads_max] | min }}"

```

(continues on next page)

(continued from previous page)

```

neutron_agent_down_time: 120
neutron_agent_polling_interval: 5
neutron_report_interval: "{{ neutron_agent_down_time | int / 2 | int }}"

neutron_dns_domain: "{{ dhcp_domain | default('openstacklocal.') }}"

# If ``neutron_num_sync_threads`` is unset, the system will use the value of
# neutron_api_threads in templates/dhcp_agent.ini.j2 for num_sync_threads.
# neutron_num_sync_threads: 4

###
### DNSMasq configuration
###
# Dnsmasq doesn't work with config_template override, a deployer
# should instead configure its own neutron_dhcp_config key/values like_
→this:
#neutron_dhcp_config:
# dhcp-option-force: "26,1500"
neutron_dhcp_config: {}

# Disable dnsmasq to resolve DNS via local resolv.conf.
# When dnsmasq_dns_servers are not set,
# and neutron_dnsmasq_noresolv is set to True, dnsmasq will reply with
# empty response on DNS requests.
neutron_dnsmasq_noresolv: False

###
### Tunable Overrides (Sorted alphabetically)
###

# These variables facilitate adding config file entries
# for anything supported by the service. See the section
# 'Overriding OpenStack configuration defaults' in the
# 'Advanced configuration' appendix of the Deploy Guide.
neutron_api_paste_ini_overrides: {}
_neutron_api_paste_ini_overrides:
  "composite:neutronapi_v2_0":
    noauth: "cors http_proxy_to_wsgi request_id catch_errors osprofiler_
→extensions neutronapiapp_v2_0"
    keystone: "cors http_proxy_to_wsgi request_id catch_errors osprofiler_
→authtoken keystonecontext extensions neutronapiapp_v2_0"

neutron_bgp_draagent_ini_overrides: {}
neutron_bgp_draagent_init_overrides: {}
neutron_calico_dhcp_agent_ini_overrides: {}
neutron_calico_dhcp_agent_init_overrides: {}
neutron_calico_felix_ini_overrides: {}
neutron_calico_felix_init_overrides: {}
neutron_dhcp_agent_ini_overrides: {}
neutron_dhcp_agent_init_overrides: {}
neutron_l3_agent_ini_overrides: {}
neutron_l3_agent_init_overrides: {}
neutron_linuxbridge_agent_ini_overrides: {}
neutron_linuxbridge_agent_init_overrides: {}
neutron_metadata_agent_ini_overrides: {}
neutron_metadata_agent_init_overrides: {}

```

(continues on next page)

(continued from previous page)

```

neutron_metering_agent_ini_overrides: {}
neutron_metering_agent_init_overrides: {}
neutron_ml2_conf_ini_overrides: {}
neutron_ml2_conf_genericswitch_ini_overrides: {}
neutron_neutron_conf_overrides: {}
neutron_nuage_conf_ini_overrides: {}
neutron_opendaylight_conf_ini_overrides: {}
neutron_openvswitch_agent_ini_overrides: {}
neutron_openvswitch_agent_init_overrides: {}
neutron_nsx_conf_ini_overrides: {}
# Provide a list of access controls to update the default policy.json with.
# These changes will be merged
# with the access controls in the default policy.json. E.g.
#neutron_policy_overrides:
# "create_subnet": "rule:admin_or_network_owner"
# "get_subnet": "rule:admin_or_owner or rule:shared"
neutron_policy_overrides: {}
_neutron_rootwrap_conf_overrides:
  DEFAULT:
    filters_path: "{{ neutron_conf_dir }}/rootwrap.d,/usr/share/neutron/
↪rootwrap"
    exec_dirs: "{{ neutron_bin }},/sbin,/usr/sbin,/bin,/usr/bin,/usr/local/
↪bin,/usr/local/sbin"
neutron_rootwrap_conf_overrides: {}

neutron_server_init_overrides: {}
neutron_sriov_nic_agent_ini_overrides: {}
neutron_sriov_nic_agent_init_overrides: {}
neutron_vpn_agent_init_overrides: {}
neutron_vpnaas_agent_ini_overrides: {}
neutron_ovn_metadata_agent_ini_overrides: {}
neutron_ovn_metadata_agent_init_overrides: {}

###
### Quotas
###

neutron_default_quota: -1
neutron_quota_floatingip: 50
neutron_quota_health_monitor: -1
neutron_quota_member: -1
neutron_quota_network: 100
neutron_quota_network_gateway: 5
neutron_quota_packet_filter: 100
neutron_quota_pool: 10
neutron_quota_port: 500
neutron_quota_router: 10
neutron_quota_security_group: 10
neutron_quota_security_group_rule: 100
neutron_quota_subnet: 100
neutron_quota_vip: 10
neutron_quota_firewall: 10
neutron_quota_firewall_policy: 10
neutron_quota_firewall_rule: 100

###

```

(continues on next page)

(continued from previous page)

```

### DB (Galera) integration
###

neutron_db_setup_host: "{{ openstack_db_setup_host | default('localhost') }}
→}"
neutron_db_setup_python_interpreter: "{{ openstack_db_setup_python_
→interpreter | default((neutron_db_setup_host == 'localhost') |
→ternary(ansible_playbook_python, ansible_facts['python']['executable']))
→}"
neutron_galera_address: "{{ galera_address | default('127.0.0.1') }}"
neutron_galera_user: neutron
neutron_galera_database: neutron
neutron_db_max_overflow: 20
neutron_db_pool_size: 120
neutron_db_pool_timeout: 30
neutron_galera_use_ssl: "{{ galera_use_ssl | default(False) }}"
neutron_galera_ssl_ca_cert: "{{ galera_ssl_ca_cert | default('/etc/ssl/
→certs/galera-ca.pem') }}"
neutron_galera_port: "{{ galera_port | default('3306') }}"

###
### Oslo Messaging
###

# RabbitMQ

neutron_oslmsg_heartbeat_in_pthread: "{{ oslmsg_heartbeat_in_pthread |
→default(False) }}"

# RPC

neutron_oslmsg_rpc_host_group: "{{ oslmsg_rpc_host_group | default(
→'rabbitmq_all') }}"
neutron_oslmsg_rpc_setup_host: "{{ (neutron_oslmsg_rpc_host_group in
→groups) | ternary(groups[neutron_oslmsg_rpc_host_group][0], 'localhost
→') }}"
neutron_oslmsg_rpc_transport: "{{ oslmsg_rpc_transport | default('rabbit
→') }}"
neutron_oslmsg_rpc_servers: "{{ oslmsg_rpc_servers | default('127.0.0.1
→') }}"
neutron_oslmsg_rpc_port: "{{ oslmsg_rpc_port | default('5672') }}"
neutron_oslmsg_rpc_use_ssl: "{{ oslmsg_rpc_use_ssl | default(False) }}"
neutron_oslmsg_rpc_userid: neutron
neutron_oslmsg_rpc_vhost: /neutron
neutron_oslmsg_rpc_ssl_version: "{{ oslmsg_rpc_ssl_version | default(
→'TLSv1_2') }}"
neutron_oslmsg_rpc_ssl_ca_file: "{{ oslmsg_rpc_ssl_ca_file | default('')
→ }}"

# Notify

neutron_oslmsg_notify_host_group: "{{ oslmsg_notify_host_group | default(
→'rabbitmq_all') }}"
neutron_oslmsg_notify_setup_host: "{{ (neutron_oslmsg_notify_host_group
→in groups) | ternary(groups[neutron_oslmsg_notify_host_group][0],
→'localhost') }}"

```

(continues on next page)

(continued from previous page)

```

neutron_oslmsg_notify_transport: "{{ oslmsg_notify_transport | default(
  ↳'rabbit') }}"
neutron_oslmsg_notify_servers: "{{ oslmsg_notify_servers | default('127.
  ↳0.0.1') }}"
neutron_oslmsg_notify_port: "{{ oslmsg_notify_port | default('5672') }}"
neutron_oslmsg_notify_use_ssl: "{{ oslmsg_notify_use_ssl |
  ↳default(False) }}"
neutron_oslmsg_notify_userid: "{{ neutron_oslmsg_rpc_userid }}"
neutron_oslmsg_notify_password: "{{ neutron_oslmsg_rpc_password }}"
neutron_oslmsg_notify_vhost: "{{ neutron_oslmsg_rpc_vhost }}"
neutron_oslmsg_notify_ssl_version: "{{ oslmsg_notify_ssl_version |
  ↳default('TLSv1_2') }}"
neutron_oslmsg_notify_ssl_ca_file: "{{ oslmsg_notify_ssl_ca_file |
  ↳default('') }}"

###
### (Qdrouterd) integration
###
# TODO(evrardjp): Change structure when more backends will be supported
neutron_oslmsg_amqp1_enabled: "{{ neutron_oslmsg_rpc_transport == 'amqp'
  ↳ }}"

###
### (RabbitMQ) integration
###

neutron_rpc_thread_pool_size: 64
neutron_rpc_conn_pool_size: 30
neutron_rpc_response_timeout: 60
neutron_rpc_workers_max: 16
neutron_rpc_workers: "{{ [(ansible_facts['processor_vcpus']//ansible_
  ↳facts['processor_threads_per_core'])|default(1), 1] | max * 2, neutron_
  ↳rpc_workers_max] | min }}"

###
### Identity (Keystone) integration
###

neutron_service_project_name: service
neutron_service_project_domain_id: default
neutron_service_user_domain_id: default
neutron_service_role_name: admin
neutron_service_user_name: neutron
neutron_service_name: neutron
neutron_service_type: network
neutron_service_description: "OpenStack Networking"
neutron_api_bind_address: "{{ openstack_service_bind_address | default('0.
  ↳0.0.0') }}"
neutron_service_port: 9696
neutron_service_proto: http
neutron_service_publicuri_proto: "{{ openstack_service_publicuri_proto |
  ↳default(neutron_service_proto) }}"
neutron_service_adminuri_proto: "{{ openstack_service_adminuri_proto |
  ↳default(neutron_service_proto) }}"
neutron_service_internaluri_proto: "{{ openstack_service_internaluri_proto
  ↳| default(neutron_service_proto) }}"

```

(continues on next page)

(continued from previous page)

```

neutron_service_publicuri: "{{ neutron_service_publicuri_proto }}://{{
  ↳external_lb_vip_address }}:{{ neutron_service_port }}"
neutron_service_publicurl: "{{ neutron_service_publicuri }}"
neutron_service_adminuri: "{{ neutron_service_adminuri_proto }}://{{
  ↳internal_lb_vip_address }}:{{ neutron_service_port }}"
neutron_service_adminurl: "{{ neutron_service_adminuri }}"
neutron_service_internaluri: "{{ neutron_service_internaluri_proto }}://{{
  ↳internal_lb_vip_address }}:{{ neutron_service_port }}"
neutron_service_internalurl: "{{ neutron_service_internaluri }}"
neutron_service_region: "{{ service_region | default('RegionOne') }}"
neutron_keystone_auth_plugin: "{{ neutron_keystone_auth_type }}"
neutron_keystone_auth_type: password
neutron_service_in_ldap: "{{ service_ldap_backend_enabled | default(False)
  ↳ }}"

###
### Telemetry integration
###

neutron_ceilometer_enabled: "{{ (groups['ceilometer_all'] is defined) and
  ↳(groups['ceilometer_all'] | length > 0) }}"

###
### Designate integration
###

neutron_designate_enabled: "{{ (groups['designate_all'] is defined) and
  ↳(groups['designate_all'] | length > 0) }}"
neutron_allow_reverse_dns_lookup: True
neutron_ipv4_ptr_zone_prefix_size: 24
neutron_ipv6_ptr_zone_prefix_size: 116

# Notifications topic for designate
neutron_notifications_designate: notifications_designate

###
### Plugins Loading
###

# Other plugins can be added to the system by simply extending the list
  ↳`neutron_plugin_base`.
# neutron_plugin_base:
#   - router
#   - firewall/firewall_v2 either one or the other, not both
#   - neutron_dynamic_routing.services.bgp.bgp_plugin.BgpPlugin
#   - vpnaas
#   - metering
#   - qos
#   - dns/subnet_dns_publish_fixed_ip either one or the other, not both
#   - port_forwarding
neutron_plugin_base:
  - router
  - metering

###
### Memcache override

```

(continues on next page)

(continued from previous page)

```
###
neutron_memcached_servers: "{{ memcached_servers }}"

###
### ML2 Plugin Configuration
###

# The neutron core plugin (ML2) is defined with neutron_plugin_type,
# you can not load multiple ML2 plugins as core.
neutron_plugin_type: 'ml2.lxb'

# Additional ML2 plugins can be loaded with neutron_plugin_types (as list)
neutron_plugin_types: []

# ml2 network type drivers to load
neutron_ml2_drivers_type: "flat,vlan,vxlan,local"

# Enable or disable L2 Population.
# When using ovs dvr it must be enabled
neutron_l2_population: "{{ neutron_plugin_type == 'ml2.ovs.dvr' }}"

neutron_vxlan_enabled: true

## The neutron multicast group address. This should be set as a host_
→variable if used.
neutron_vxlan_group: "239.1.1.1"

# The neutron multicast time-to-live. Number of L3 hops before routers_
→will drop the traffic
neutron_vxlan_ttl: 32

neutron_sriov_excluded_devices: ""

# neutron_local_ip is used for the VXLAN local tunnel endpoint
neutron_local_ip: 127.0.0.1

## Set this variable to configure the provider networks that will be_
→available
## When setting up networking in things like the ml2_conf.ini file._
→Normally
## this will be defined as a host variable used within neutron as network_
→configuration
## are likely to differ in between hosts.
# neutron_provider_networks:
#   network_flat_networks: "flat"
#   network_mappings: "flat:eth12,vlan:eth11"
#   network_types: "vxlan,flat,vlan"
#   network_vlan_ranges: "vlan:1:1,vlan:1024:1025"
#   network_vxlan_ranges: "1:1000"
#   network_geneve_ranges: "1:1000"
#   network_sriov_mappings: "vlan:p4p1"

###
### L3 Agent Plugin Configuration
###
```

(continues on next page)

(continued from previous page)

```
# L3HA configuration options
neutron_ha_vrrp_auth_type: PASS
neutron_l3_ha_net_cidr: 169.254.192.0/18

neutron_l3_cleanup_on_shutdown: False

###
### DHCP Agent Plugin Configuration
###

# Comma-separated list of DNS servers which will be used by dnsmasq as
# forwarders.
neutron_dnsmasq_dns_servers: ""

# Limit number of leases to prevent a denial-of-service.
neutron_dnsmasq_lease_max: 16777216

# Specify if dnsmasq should send a route to metadata server through DHCP
# 121 message to VM
neutron_dnsmasq_force_metadata: False

###
### Metadata Agent Plugin Configuration
###

# If ``neutron_metadata_workers`` is unset the system will use half the
# number of available VCPUs to
# compute the number of api workers to use with a default capping value of
# 16.
neutron_metadata_workers: 16
neutron_metadata_backlog: 4096

# The port used by neutron to access the nova metadata service.
neutron_nova_metadata_port: "{{ nova_metadata_port | default(8775) }}"

# The protocol used by neutron to access the nova metadata service.
neutron_nova_metadata_protocol: "{{ nova_metadata_protocol | default('http') }}"

# If the nova_metadata_protocol is using a self-signed cert, then
# this flag should be set to a boolean True.
neutron_nova_metadata_insecure: "{{ nova_metadata_insecure |
    default(False) }}"

###
### VPNaaS Configuration
###

# See VPNaaS documentation for driver/service provider selection
# in case you want to override it.
neutron_driver_vpnaas: "{{ _neutron_driver_vpnaas }}"
neutron_vpnaas_service_provider: "{{ _neutron_vpnaas_service_provider }}"

# Calico Felix agent upstream settings
calico_felix_url: "https://github.com/projectcalico/felix/releases/
    download/{{ calico_felix_version }}/calico-felix-amd64"
```

(continues on next page)

(continued from previous page)

```
calico_felix_version: v3.7.0
calico_felix_sha256: ↵
↪ae0bed304702097cee0ad5d9b4abb07b263deeb46ac21f2bcb0118d5bf439f46
calico_felix_validate_certs: yes

# OVN Defaults
neutron_ovn_primary_cluster_node: "{{ groups[neutron_services['neutron-ovn-
↪northd']]['group']] | first }}"
neutron_ovn_northd_service_name: ovn-northd
neutron_ovn_controller_service_name: ovn-controller
neutron_ovn_l3_scheduler: leastloaded
neutron_ovn_ip: "{{ internal_lb_vip_address }}"
neutron_ovsdb_manager: tcp:6640:127.0.0.1

###
### DPDK Configuration
###

ovs_datapath: "netdev"
ovs_dpdk_pci_addresses: []
ovs_dpdk_driver: vfio-pci
ovs_dpdk_support: False
ovs_dpdk_lcore_mask: 1
ovs_dpdk_pmd_cpu_mask: 2
ovs_dpdk_socket_mem: "1024"
ovs_dpdk_nr_1g_pages: 0
ovs_dpdk_nr_2m_pages: 0
```

**CHAPTER
FOURTEEN**

DEPENDENCIES

This role needs pip ≥ 7.1 installed on the target host.

EXAMPLE PLAYBOOK

```
---
- name: Installation and setup of Neutron
  hosts: neutron_all
  user: root
  roles:
    - { role: "os_neutron", tags: [ "neutron-install", "neutron-config" ] }
  vars:
    neutron_galera_address: "{{ internal_lb_vip_address }}"
    galera_root_user: root
  vars_prompt:
    - name: "galera_root_password"
      prompt: "What is galera_root_password?"
```

**CHAPTER
SIXTEEN**

TAGS

This role supports two tags: `neutron-install` and `neutron-config`. The `neutron-install` tag can be used to install and upgrade. The `neutron-config` tag can be used to maintain the configuration of the service.