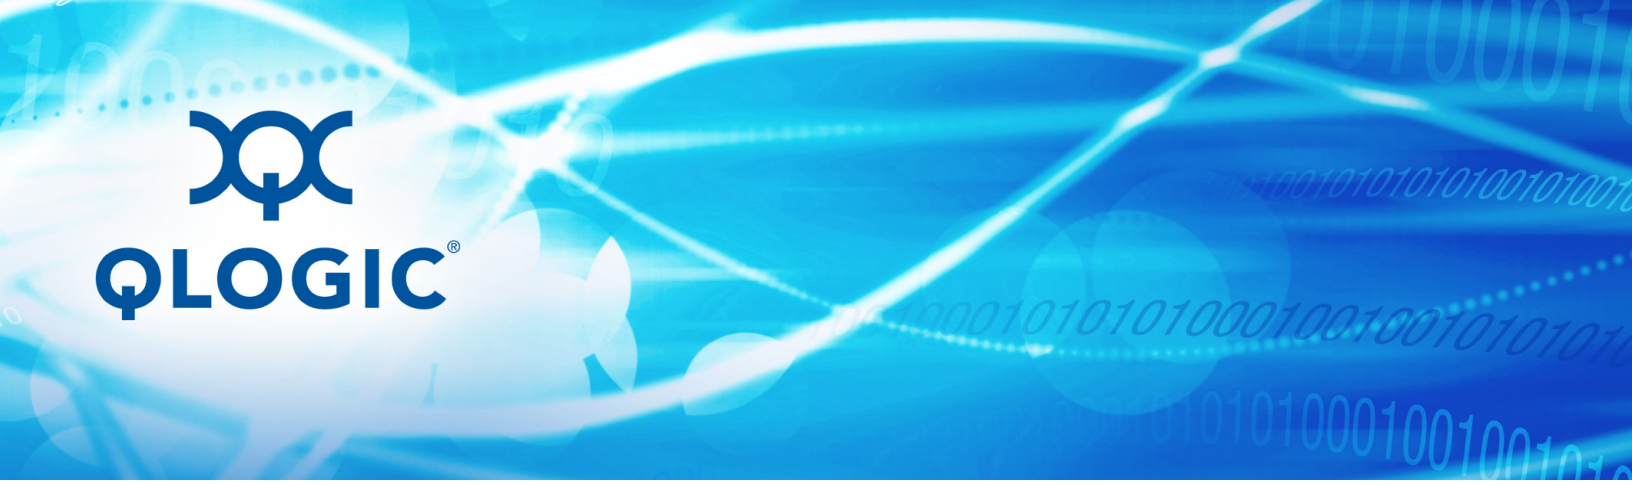

User's Guide

OpenStack Deployment with SR-IOV Configuration

QLogic 3400, 8400, and 45000 Series Ethernet Adapters

83840-546-02 B

Document Revision History

Rev 01, May 15, 2015

Rev 02, May 22, 2015

Rev 03, March 22, 2016

Rev A, April 19, 2016

Rev B, August 29, 2016

Changes	Sections Affected
Added Liberty and Mitaka to the applicable OpenStack platforms. Added a footnote associating 3400/8400/45000 Series Adapters with the 57800 Series Adapters and QL4521*x* adapters.	"Intended Audience" on page vii
Added Liberty and Mitaka to the applicable OpenStack platforms. Added a NOTE indicating that support for the 45000 Series is limited to the QL4521*x* adapters.	"Audience" on page 1
In the first bullet, added a NOTE indicating that support for the 45000 Series is limited to the QL4521*x* adapters. In the second bullet, added links to information about OpenStackinstallation for Linux for Liberty and Mitaka.	"Prerequisites" on page 3
Added new steps 8, 9, and 10.	"Enabling SR-IOV on the 3400/8400/45000 Series Adapters (Compute Node)" on page 6
Added new Steps 3 and 4. In Step 6, added a third command: `# Systemctl restart open-stack-nova-scheduler.service`	"Configuring SR-IOV on the Controller Node " on page 7

In Step 3, added a qualifier for Juno and Kilo to the first command and added a new command for Liberty and Mitaka. In Step 4, added a qualifier for Juno and Kilo to the `ml2_conf.ini` file edits and added `sriov_agent.ini` file edits for Liberty and Mitaka. In Steps 5 and 6, added a qualifier for Juno and Kilo only. In Step 7, added a qualifier for Juno and Kilo to the first command and added new commands for Liberty and Mitaka.	"Configuring SR-IOV on the Compute Node" on page 9

Table of Contents

Preface

Intended Audience

This guide is intended for OpenStack® (Juno, Kilo, Liberty, and Mitaka) users who want to configure single root input/output virtualization (SR-IOV) with the QLogic® 3400/8400/45000 Series Adapters[1].

What Is in This Guide

This preface specifies the intended audience, explains the typographic conventions used in this guide, lists related documents, and provides technical support and contact information.

The remainder of this guide is organized into the following chapters:

- Chapter 1 Introduction describes the audience for this guide and outlines the OpenStack architecture.
- Chapter 2 Example Architecture describes the OpenStack 3 node architecture.
- Chapter 3 Prerequisites describes what you need to use OpenStack with SR-IOV configuration using the QLogic 3400/8400/45000 Series Adapters.
- Chapter 4 SR-IOV Configuration Using the 3400/8400/45000 Series Adapters describes how to configure SR-IOV using the QLogic 3400/8400/45000 Series Adapters.
- Chapter 5 Configuration Deployment with SR-IOV describes how to deploy OpenStack with SR-IOV.
- Chapter 6 Testing provides a sample testing procedure.
- Chapter 7 Known Issues describes a known issue with open virtual search (OVS) and OpenStack.

Related Materials

For information about downloading documentation from the QLogic Web site, see “Downloading Updates” on page x.

[1] Includes 578*xx*-based Adapters and FastLinQ™ QL4521*x* 25-gigabit Ethernet (GbE) Adapters

Documentation Conventions

This guide uses the following documentation conventions:

- Text in blue font indicates a hyperlink (jump) to a figure, table, or section in this guide, and links to Web sites are shown in underlined blue. For example:
  - Table 9-2 lists problems related to the user interface and remote agent.
  - See "Installation Checklist" on page 6.
  - For more information, visit www.qlogic.com.
- Text in **bold** font indicates user interface elements such as a menu items, buttons, check boxes, or column headings. For example:
  - Click the **Start** button, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.
  - Under **Notification Options**, select the **Warning Alarms** check box.
- Text in `Courier` font indicates a file name, directory path, or command line text. For example:
  - To return to the root directory from anywhere in the file structure: Type `cd /root` and press ENTER.
  - Enter the following command: `sh ./install.bin`
- Key names and key strokes are indicated with UPPERCASE:
  - Press CTRL+P.
  - Press the UP ARROW key.
- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:
  - For a complete listing of license agreements, refer to the *QLogic Software End User License Agreement*.
  - What are *shortcut keys*?
  - To enter the date type *mm/dd/yyyy* (where *mm* is the month, *dd* is the day, and *yyyy* is the year).
- Topic titles between quotation marks identify related topics within this manual.

- Command syntax conventions include the following:
  - Plain text indicates items that you must type as shown. For example:
    - `qaucli -pr nic -ei`
  - < > (angle brackets) indicate a variable whose value you must specify. For example:
    - `<serial_number>`
  - [ ] (square brackets) indicate an optional parameter. For example:
    - `[<file_name>]` means specify a file name, or omit it to select the default file name.
  - | (vertical bar) indicates mutually exclusive options; select one option only. For example:
    - `on|off`
    - `1|2|3|4`
  - ... (ellipsis) indicates that the preceding item may be repeated. For example:
    - `x...` means *one* or more instances of `x`.
    - `[x...]` means *zero* or more instances of `x`.
  - ( ) (parentheses) and { } (braces) are used to avoid logical ambiguity. For example:
    - `a|b c` is ambiguous
      `{(a|b) c}` means `a` or `b`, followed by `c`
      `{a|(b c)}` means either `a,` or `b c`

License Agreements

Refer to the *QLogic Software End User License Agreement* for a complete listing of all license agreements affecting the QLogic 3400/8400/45000 Series Adapters.

Technical Support

Customers should contact their authorized maintenance provider for technical support of their QLogic products. QLogic-direct customers may contact QLogic Technical Support; others will be redirected to their authorized maintenance provider. Visit the QLogic support Web site listed in Contact Information for the latest firmware and software updates.

For details about available service plans, or for information about renewing and extending your service, visit the Service Program Web page at www.qlogic.com/Support/Pages/ServicePrograms.aspx.

Downloading Updates

The QLogic Web site provides periodic updates to product firmware, software, and documentation.

To download firmware, software, and documentation:

1. Go to the QLogic Downloads and Documentation page: driverdownloads.qlogic.com.
2. Type the QLogic model name in the search box.
3. In the search results list, locate and select the firmware, software, or documentation for your product.
4. View the product details Web page to ensure that you have the correct firmware, software, or documentation. For additional information, click **Read Me** and **Release Notes** under Support Files.
5. Click **Download Now**.
6. Save the file to your computer.
7. If you have downloaded firmware, software, drivers, or boot code, follow the installation instructions in the *Readme* file.

Instead of typing a model name in the search box, you can perform a guided search as follows:

1. Click the product type tab: **Adapters**, **Switches**, **Routers**, or **ASICs**.
2. Click the corresponding button to search by model or operating system.
3. Click an item in each selection column to define the search, and then click **Go**.
4. Locate the firmware, software, or document you need, and then click the item's name or icon to download or open the item.

Training

QLogic Global Training maintains a Web site at www.qlogictraining.com offering online and instructor-led training for all QLogic products. In addition, sales and technical professionals may obtain Associate and Specialist-level certifications to qualify for additional benefits from QLogic.

Contact Information

QLogic Technical Support for products under warranty is available during local standard working hours excluding QLogic Observed Holidays. For customers with extended service, consult your plan for available hours. For Support phone numbers, see the Contact Support link at support.qlogic.com.

Support Headquarters	QLogic Corporation 12701 Whitewater Drive Minnetonka, MN 55343 USA
QLogic Web Site	www.qlogic.com
Technical Support Web Site	support.qlogic.com
Technical Support E-mail	support@qlogic.com
Technical Training E-mail	training@qlogic.com

Knowledge Database

The QLogic knowledge database is an extensive collection of QLogic product information that you can search for specific solutions. QLogic is constantly adding to the collection of information in the database to provide answers to your most urgent questions. Access the database from the QLogic Support Center: support.qlogic.com.

1 Introduction

This chapter describes the audience for this guide and outlines the OpenStack architecture.

Audience

This guide is intended for OpenStack (Juno, Kilo, Liberty, and Mitaka) users who want to configure single root input/output virtualization (SR-IOV) with the QLogic 3400/8400/45000 Series Adapters.

NOTE

At the time of publication, support for the 45000 Series Adapters is limited to the FastLinQ™ QL4521*x* 25GbE Adapters.

What is OpenStack?

The OpenStack project is an open source cloud computing platform that supports all types of cloud environments and works as infrastructure as a service (IAAS).

The technology consists of a series of interrelated projects that control pools of processing, storage, and networking resources throughout a data center, which users manage through a web-based dashboard, command line tools, or representational state transfer (REST) APIs.

OpenStack (starting with Juno) adds inbox support to request virtual machine (VM) access to the virtual network through SR-IOV NIC. With the introduction of SR-IOV-based NICs, each SR-IOV port is associated with virtual functions (VFs). SR-IOV ports are provided by hardware-based virtual Ethernet bridging (HW VEB).

2 Example Architecture

This chapter describes the OpenStack 3 node architecture.

OpenStack Nodes

This guide is intended for use with OpenStack 3 node architecture, which consists of the following nodes:

- Controller
- Compute
- Network

These nodes are described in the following sections.

Controller Node

In the example architecture used in this guide, the Controller node runs the Identity service (Keystone), Image Service (Glance), management portions of the Compute service (Nova Management) and Networking service (Neutron Server/modular layer 2 (ML2) plug-in), networking plug-in, and the dashboard (Horizon). The architecture also includes supporting services such as a database (mysql), message broker (Rabbitmq), and network time protocol (NTP).

Compute Node

The Compute node runs the hypervisor portion of Compute, which operates tenant virtual machines.

In this architecture, the Compute node uses a kernel-based virtual machine (KVM) as the hypervisor (KVM is the default hypervisor). The Compute node runs the Networking node plug-in (ML2), layer 2 agent open virtual search (OVS), and NIC switch agent (SR-IOV switch).

Network Node

The Network node runs the Networking plug-in (ML2), layer 3 agent, and DHCP agent. This node also handles external (Internet) connectivity for tenant VMs or the Compute node.

3 Prerequisites

This chapter describes what you need to use OpenStack with SR-IOV configuration using the QLogic 3400/8400/45000 Series Adapters.

You need the following knowledge and equipment:

- One or more QLogic 3400/8400/45000 Series Adapters

  NOTE

  At the time of publication, support for the 45000 Series Adapters is limited to the FastLinQ QL4521*x* 25GbE Adapters.

- An understanding of OpenStack deployment and experience with OpenStack 3 node configuration using flat mode networking with virtual LAN (VLAN) with the ML2/OVS plug-in for public and private (that is, net1) networks.

  More information about OpenStack 3 node configuration on Red Hat® 7 can be found at the following locations:

  OpenStack (Juno):
  http://docs.openstack.org/juno/install-guide/install/yum/content/

  OpenStack (Kilo):
  http://docs.openstack.org/kilo/install-guide/install/yum/content/

  Open Stack (Liberty)[1]
  http://docs.openstack.org/liberty/install-guide-rdo/

  Open Stack (Mitaka)[1]
  http://docs.openstack.org/mitaka/install-guide-rdo/

- The Compute node must have one or more QLogic 3400/8400/45000 Series Adapters present in system with SR-IOV support enabled.

[1] The OpenStack configuration described in this link uses VXLAN and Linux Bridge agent. However, the QLogic SR-IOV configuration uses and has been tested only with OpenStack 3 node configuration with VLAN and open vSwitch (OVS). Furthermore, be sure to configure the OpenStack 3 node configuration with kernel-based virtual machine (KVM).

- For the VM instance, QLogic uses a Red Hat 7 KVM Guest image to test SR-IOV configurations that support the 3400/8400/45000 Series Adapters' inbox or out-of-box driver, as described here: http://docs.openstack.org/image-guide/content/ch_obtaining_images.html
- OpenStack 3 node configuration without SR-IOV support must work with flat networking/VLAN. You must be able to assign floating IP access and log in to the VM.

4 SR-IOV Configuration Using the 3400/8400/45000 Series Adapters

This chapter describes how to configure SR-IOV using the QLogic 3400/8400/45000 Series Adapters, as described in the following sections:

- Enabling SR-IOV on the 3400/8400/45000 Series Adapters (Compute Node)
- Configuring SR-IOV on the Controller Node
- Configuring SR-IOV on the Compute Node

NOTE

The steps in the following sections can easily be accomplished using the QLogic provided OpenStack Utility, available for download from: http://www.qlogic.com/solutions/Pages/OpenStack.aspx. Two packages are provided:

- **qlgcOpenStack.zip**. Standalone package containing source files
- **qlgcOpenStackEc.zip**. Standalone package containing source files and Eclipse project files to simplify the customization of the QLogic scripts using Eclipse

Enabling SR-IOV on the 3400/8400/45000 Series Adapters (Compute Node)

To enable SR-IOV on the 3400/8400/45000 Series Adapters in the Compute node:

1. Ensure that one or more 3400/8400/45000 Series Adapters with SR-IOV support are present on the Compute node.
2. Install the appropriate NIC driver on the Compute node:
   - ❑ bnx2x (3400 and 8400 Series Adapters)
   - ❑ qede (45000 Series Adapters)
3. Load the NIC drivers by issuing the following command:

   `modprobe bnx2x` (3400 and 8400 Series Adapters)

   `modprobe qede` (45000 Series Adapters; this command installs both the `qede` and `qed` driver)
4. Ensure that the port is up by issuing the following command:

   `ifup ethx / ifconfig ethx up`
5. To create eight VFs, issue the following command:

   `echo 8 > /sys/bus/pci/devices/0000\:0x\:00.x/sriov_numvfs`
6. Verify that the virtual interfaces are visible to host OS by issuing the following command:

   `lspci |grep Qlogic`
7. Enable SR-IOV on the port where the tenant network has been defined.
8. Modify the `/boot/grub2/grub.cfg` file to enable `intel-iommu` on the compute node.

   ```
   linux16 /vmlinuz-3.10.0-327.22.2.el7.x86_64
   root=/dev/mapper/centos-root ro crashkernel=auto
   rd.lvm.lv=centos/root rd.lvm.lv=centos/swap rhgb quiet
   intel_iommu=on LANG=en_US.UTF-8
   initrd16 /initramfs-3.10.0-327.22.2.el7.x86_64.img
   ```
9. Reboot the host.

10. Add `intel_iommu=on` to the command line of the kernel from which you will boot. Type the following command:

```
cat /proc/cmdline show intel_iommu=on
BOOT_IMAGE=/vmlinuz-3.10.0-327.22.2.el7.x86_64
root=/dev/mapper/centos-root ro crashkernel=auto
rd.lvm.lv=centos/root rd.lvm.lv=centos/swap rhgb quiet
intel_iommu=on LANG=en_US.UTF-8
```

Configuring SR-IOV on the Controller Node

Since the Neutron server is running on the Controller node, the file configuration for the Neutron server must be changed to enable SR-IOV, as described in the following procedure.

To configure SR-IOV on the Controller node:

1. Modify the ML2 Neutron plug-in by editing the `/etc/neutron/plugins/ml2/ml2_conf.ini` file on Controller node as follows:

```
[ml2]
tenant_network_types = flat,vlan
type_drivers = vlan
mechanism_drivers = openvswitch,sriovnicswitch

[ml2_type_vlan]
network_vlan_ranges = physnet2:1000:2000,physnet3:2001:2100

[securitygroup]
firewall_driver = neutron.agent.firewall.NoopFirewallDriver
```

2. Add the supported PCI vendor VF devices defined by `vendor_id:product_id` according to the PCI ID Repository by editing the `/etc/neutron/plugins/ml2/ml2_conf_sriov.ini` file as follows:

```
[ml2_sriov]
supported_pci_vendor_devs = 14e4:16xx (3400 and 8400 Series
Adapters)
supported_pci_vendor_devs = 1077:16xx (45000 Series Adapters)
```

3. Modify the `ml2_conf_sriov.ini` file to add the following statement:

```
[ml2_sriov]
Agent_required=True
```

4. On the controller node running nova-scheduler, modify the `/etc/nova/nova.conf` file [DEFAULT] section to add the following:

   - `PciPassthroughFilter` to the `scheduler_default_filters` parameter
   - A new line for the `scheduler_available_filters` parameter:

```
[DEFAULT]
scheduler_default_filters =
RetryFilter, AvailabilityZoneFilter, RamFilter,
ComputeFilter, ComputeCapabilitiesFilter,
ImagePropertiesFilter, ServerGroupAntiAffinityFilter,
ServerGroupAffinityFilter, PciPassthroughFilter

scheduler_available_filters =
nova.scheduler.filters.all_filters
```

5. Change the Neutron server configuration.

   The Neutron server must be run with the following two configuration files:

   - `/etc/neutron/plugins/ml2/ml2_conf.in`
   - `/etc/neutron/plugins/ml2/ml2_conf_sriov.ini`

   a. Navigate to the configuration files' location by issuing the following command:

```
# cd /usr/lib/systemd/system
```

   b. Edit the `neutron-server.service` file and add the `ml2_conf_sriov.ini` configuration file as follows:

```
ExecStart=/usr/bin/neutron-server --config-file
/usr/share/neutron/neutron-dist.conf --config-file
/etc/neutron/neutron.conf --config-file
/etc/neutron/plugin.ini --config-file
/etc/neutron/plugins/ml2/ml2_conf_sriov.ini --log-file
/var/log/neutron/server.log
```

6. Restart the Neutron server by issuing the following commands:

```
# systemctl daemon-reload
# systemctl restart neutron-server.service
# systemctl restart openstack-nova-scheduler.service
```

Configuring SR-IOV on the Compute Node

To configure SR-IOV on the Compute node, associate the available VF with the physical network as follows:

1. Add `pci_passthrough_whitelist` to the `/etc/nova/nova.conf` file as follows:

   ```
   pci_passthrough_whitelist={"devname": "p5p2",
   "physical_network":"physnet3"}
   ```

2. Restart the OpenStack Nova service by issuing the following command:

   ```
   # systemctl restart openstack-nova-compute.service
   ```

3. Install the SR-IOV Neutron agent on the Compute node by issuing the following command:

   - ❑ For Juno and Kilo:

     ```
     # yum install neutron-sriov-nic-agent
     ```

   - ❑ For Liberty and Mitaka:

     ```
     # yum install openstack-neutron-sriov-nic-agent.noarch
     ```

4. Configure ML2 on the Compute node.

   - ❑ For Juno and Kilo, edit the node's `ml2_conf.ini` file to add the `NoopFirewallDriver`, as follows:

     ```
     [securitygroup]
     firewall_driver =
     neutron.agent.firewall.NoopFirewallDriver
     ```

   - ❑ For Liberty and Mitaka, modify the `sriov_agent.ini` file as follows:

     ```
     [sriov_nic]
     Physical_device_mapping = physnet3:p5p2

     [securitygroup]
     firewall_driver =
     neutron.agent.firewall.NoopFirewallDriver
     ```

5. **For Juno and Kilo only—**To run the Neutron SR-IOV NIC agent service, first navigate to the configuration file location by issuing the following command:

   ```
   # cd /usr/lib/systemd/system
   ```

6. **For Juno and Kilo only—**Edit the `neutron-sriov-nic-agent.service` file and add the `ml2_conf_sriov.ini` configuration file as follows:

```
ExecStart=/usr/bin/neutron-sriov-nic-agent --config-file
/usr/share/neutron/neutron-dist.conf --config-file
/etc/neutron/neutron.conf --config-file
/etc/neutron/plugins/ml2/ml2_conf_sriov.ini --log-file
/var/log/neutron/sriov-nic-agent.log
```

7. Restart the Neutron Open vSwitch agent service.

   - For Juno and Kilo, issue the following command:

```
# service neutron-plugin-openvswitch-agent restart
```

   - For Liberty and Mitaka, issue the following commands:

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

5 Configuration Deployment with SR-IOV

This chapter describes how to deploy OpenStack with SR-IOV, as follows:

Creating an SR-IOV Network

To create an SR-IOV network:

1. Create an SR-IOV network with a VLAN network type by issuing the following command:

   ```
   # neutron net-create --provider:physical_network=physnet3
   --provider:network_type=vlan sriov
   ```

2. Issue the following command to view the configuration for OpenStack 3 nodes:

   ```
   # [root@network1 ~]# neutron net-list
   ```

   Following is a sample output.

```
[root@network1 ~]# neutron net-list
+--------------------------------------+--------+-------------------------------------------------------+
| id                                   | name   | subnets                                               |
+--------------------------------------+--------+-------------------------------------------------------+
| 7d8886cf-e8a1-4934-8078-df63c8ad5520 | net1   | ce4699c7-690c-4391-8d6f-5fe5bb580ebe 12.0.0.0/24      |
| d434ac7a-f6be-45c8-9bcc-af9414d7bb4a | sriov  | d94e8385-c643-4f0f-80f5-3e2306995cab 11.0.0.0/24      |
| 19573474-02b3-45da-9cfb-e3d8a4b79fc0 | public | 5dfea63c-8181-4ea3-aab7-a15d4063861b 172.28.0.0/20    |
+--------------------------------------+--------+-------------------------------------------------------+
```

The values in the `name` column are described in the following paragraphs.

`net1`	A private network that provides internal network access for instances of the Compute node. Public and private networks are configured in the prerequisite section.
`sriov`	An SR-IOV network for VM access to a virtual network using SR-IOV NIC.
`public`	An external network that provides Internet access for instances of the Compute node using a network address translation (NAT) /floating IP address and a qualified security group.

Creating the Subnet “subnet2” for the SR-IOV Network

To create the subnet2 subnet for the SR-IOV network:

1. Create a subnet attached to the SR-IOV network by issuing the following command:

   ```
   # neutron subnet-create sriov --name subnet2  11.0.0.0/24
   ```

2. Issue the following command to view the subnet network:

   ```
   # neutron subnet-list
   ```

   Following is a sample output.

```
[root@network1 ~]# neutron subnet-list
+--------------------------------------+--------------+--------------+--------------------------------------------------+
| id                                   | name         | cidr         | allocation_pools                                 |
+--------------------------------------+--------------+--------------+--------------------------------------------------+
| ce4699c7-690c-4391-8d6f-5fe5bb580ebe | subnet1      | 12.0.0.0/24  | {"start": "12.0.0.2", "end": "12.0.0.254"}       |
| d94e8385-c643-4f0f-80f5-3e2306995cab | subnet2      | 11.0.0.0/24  | {"start": "11.0.0.2", "end": "11.0.0.254"}       |
| 5dfea63c-8181-4ea3-aab7-a15d4063861b | public-subnet| 172.28.0.0/20| {"start": "172.28.11.231", "end": "172.28.11.250"} |
+--------------------------------------+--------------+--------------+--------------------------------------------------+
```

The values in the `name` column are described in the following paragraphs.

`subnet1`	A private subnet that uses DHCP Private subnets are configured in the prerequisite section.
`subnet2`	An SRI-OV subnet
`public-subnet`	A subnet for external connectivity

Creating the SR-IOV Port

To create the SR-IOV port:

1. Issue the following command to create an SR-IOV port:

   ```
   # neutron port-create <sriov net id>  --name  <port_name>
   --binding:vnic-type direct --device_owner network:dhcp
   ```

   Following is a sample output.

```
[root@network1 ~]# neutron port-create d434ac7a-f6be-45c8-9bcc-af9414d7bb4a --name sriov_port9 --binding:vnic-type direct --device_owner network:dhcp
Created a new port:
+-----------------------+--------------------------------------------------------------------------------------+
| Field                 | Value                                                                                |
+-----------------------+--------------------------------------------------------------------------------------+
| admin_state_up        | True                                                                                 |
| allowed_address_pairs |                                                                                      |
| binding:host_id       |                                                                                      |
| binding:profile       | {}                                                                                   |
| binding:vif_details   | {}                                                                                   |
| binding:vif_type      | unbound                                                                              |
| binding:vnic_type     | direct                                                                               |
| device_id             |                                                                                      |
| device_owner          | network:dhcp                                                                         |
| fixed_ips             | {"subnet_id": "d94e8385-c643-4f0f-80f5-3e2306995cab", "ip_address": "11.0.0.9"}      |
| id                    | 1987b21c-f071-4f65-b1ff-509d24f306d4                                                 |
| mac_address           | fa:16:3e:ea:e0:0c                                                                    |
| name                  | sriov_port9                                                                          |
| network_id            | d434ac7a-f6be-45c8-9bcc-af9414d7bb4a                                                 |
| security_groups       |                                                                                      |
| status                | DOWN                                                                                 |
| tenant_id             | 28b7418d330c4fb08d753c84f10c4eec                                                     |
+-----------------------+--------------------------------------------------------------------------------------+
```

2. To show the port information, issue the following command:

   ```
   # neutron port-show <ID>
   ```

   Following is a sample output.

```
[root@network1 ~]# neutron port-show 1987b21c-f071-4f65-b1ff-509d24f306d4
+-----------------------+--------------------------------------------------------------------------------------+
| Field                 | Value                                                                                |
+-----------------------+--------------------------------------------------------------------------------------+
| admin_state_up        | True                                                                                 |
| allowed_address_pairs |                                                                                      |
| binding:host_id       |                                                                                      |
| binding:profile       | {}                                                                                   |
| binding:vif_details   | {}                                                                                   |
| binding:vif_type      | unbound                                                                              |
| binding:vnic_type     | direct                                                                               |
| device_id             |                                                                                      |
| device_owner          | network:dhcp                                                                         |
| extra_dhcp_opts       |                                                                                      |
| fixed_ips             | {"subnet_id": "d94e8385-c643-4f0f-80f5-3e2306995cab", "ip_address": "11.0.0.9"}      |
| id                    | 1987b21c-f071-4f65-b1ff-509d24f306d4                                                 |
| mac_address           | fa:16:3e:ea:e0:0c                                                                    |
| name                  | sriov_port9                                                                          |
| network_id            | d434ac7a-f6be-45c8-9bcc-af9414d7bb4a                                                 |
| security_groups       |                                                                                      |
| status                | DOWN                                                                                 |
| tenant_id             | 28b7418d330c4fb08d753c84f10c4eec                                                     |
+-----------------------+--------------------------------------------------------------------------------------+
```

   Make sure the output shows the correct `vif_type` and `vnic_type`.

Creating a VM with the SR-IOV Network

This section describes how to create a VM with multiple networks: a private network for a floating IP address and an SR-IOV network with direct VM access using a VF.

To create a VM with the SR-IOV network:

1. Issue the following command to create a VM:

   ```
   # nova boot --flavor m1.medium --image <image-id>  --nic
   net-id=<net id of private network net1>  --nic
   port-id=<port id of sriov network from port-create command in
   “Configuring SR-IOV on the Compute Node” on page 9>  <VM name>
   ```

   In this command, a `–nic net-id` is given for the private network (`net1`) and a `–nic port-id` is given for the SR-IOV network with a Red Hat 7 KVM guest image.

   Following is a sample output.

```
[root@network1 ~]# nova boot --flavor m1.medium --image 6292c70f-f327-4f82-be36-547afb9febcb  --nic net-id=7d8886cf-e8a1-4934-8078-df63c8ad5520 --nic port-id=1987b21c-f
071-4f65-b1ff-509d24f306d4 VNM4
+--------------------------------------+--------------------------------------+
| Property                             | Value                                |
+--------------------------------------+--------------------------------------+
| OS-DCF:diskConfig                    | MANUAL                               |
| OS-EXT-AZ:availability_zone          | nova                                 |
| OS-EXT-SRV-ATTR:host                 | -                                    |
| OS-EXT-SRV-ATTR:hypervisor_hostname  | -                                    |
| OS-EXT-SRV-ATTR:instance_name        | instance-00000072                    |
| OS-EXT-STS:power_state               | 0                                    |
| OS-EXT-STS:task_state                | scheduling                           |
| OS-EXT-STS:vm_state                  | building                             |
| OS-SRV-USG:launched_at               | -                                    |
| OS-SRV-USG:terminated_at             | -                                    |
| accessIPv4                           |                                      |
| accessIPv6                           |                                      |
| adminPass                            | 9dMK3qkVY4jb                         |
| config_drive                         |                                      |
| created                              | 2015-03-19T09:02:10Z                 |
| flavor                               | m1.medium (3)                        |
| hostId                               |                                      |
| id                                   | de88aa65-49c5-4080-9468-86d3c83f8a03 |
| image                                | rhx (6292c70f-f327-4f82-be36-547afb9febcb) |
| key_name                             | -                                    |
| metadata                             | {}                                   |
| name                                 | VNM4                                 |
| os-extended-volumes:volumes_attached | []                                   |
| progress                             | 0                                    |
| security_groups                      | default                              |
| status                               | BUILD                                |
| tenant_id                            | 28b7418d330c4fb08d753c84f10c4eec     |
| updated                              | 2015-03-19T09:02:10Z                 |
| user_id                              | 681a6563531b488a99bad094976a89f8     |
+--------------------------------------+--------------------------------------+
```

2. Issue the following command to ensure that the VM is up and running:

   ```
   # nova list
   ```

   Following is a sample output.

```
@network1 ~]# nova list
+--------------------------------------+------+--------+------------+-------------+-----------------------------------------
| Name | Status | Task State | Power State | Networks
+--------------------------------------+------+--------+------------+-------------+-----------------------------------------
81ad6-a0af-46e9-8f0b-e09cfeaa8065 | VNM1 | ACTIVE | -          | Running     | net1=12.0.0.4, 172.28.11.234; sriov=11.0.0.
5c2c2-89dd-4aa6-b9af-9a61f5d38216 | VNM2 | ACTIVE | -          | Running     | net1=12.0.0.5, 172.28.11.235; sriov=11.0.0.
b990b-cbd1-4e6e-8d9b-d60c38ce435f | VNM3 | ACTIVE | -          | Running     | net1=12.0.0.6, 172.28.11.236; sriov=11.0.0.
8aa65-49c5-4080-9468-86d3c83f8a03 | VNM4 | ACTIVE | -          | Running     | net1=12.0.0.7; sriov=11.0.0.9
+--------------------------------------+------+--------+------------+-------------+-----------------------------------------
```

3. Assign a floating IP address to the private network to access using ssh. Following is a sample output.

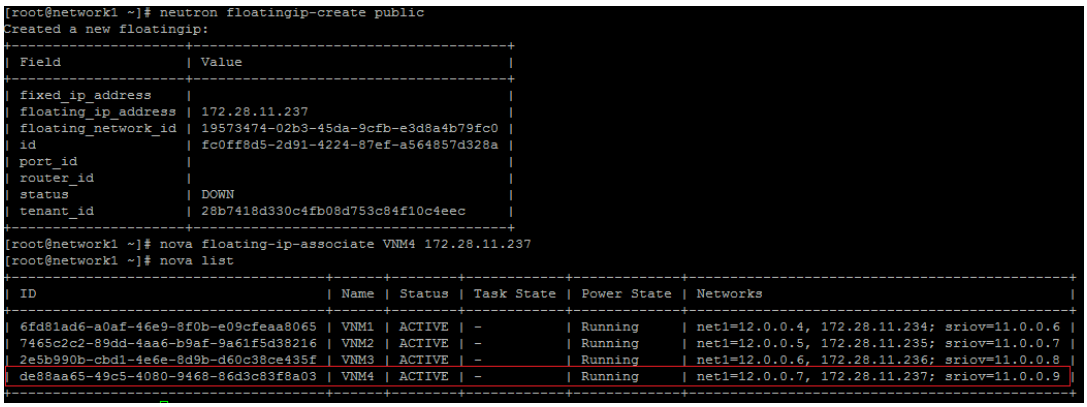

```
[root@network1 ~]# neutron floatingip-create public
Created a new floatingip:
+---------------------+--------------------------------------+
| Field               | Value                                |
+---------------------+--------------------------------------+
| fixed_ip_address    |                                      |
| floating_ip_address | 172.28.11.237                        |
| floating_network_id | 19573474-02b3-45da-9cfb-e3d8a4b79fc0 |
| id                  | fc0ff8d5-2d91-4224-87ef-a564857d328a |
| port_id             |                                      |
| router_id           |                                      |
| status              | DOWN                                 |
| tenant_id           | 28b7418d330c4fb08d753c84f10c4eec     |
+---------------------+--------------------------------------+
[root@network1 ~]# nova floating-ip-associate VNM4 172.28.11.237
[root@network1 ~]# nova list
+--------------------------------------+------+--------+------------+-------------+-----------------------------------------------+
| ID                                   | Name | Status | Task State | Power State | Networks                                      |
+--------------------------------------+------+--------+------------+-------------+-----------------------------------------------+
| 6fd81ad6-a0af-46e9-8f0b-e09cfeaa8065 | VNM1 | ACTIVE | -          | Running     | net1=12.0.0.4, 172.28.11.234; sriov=11.0.0.6 |
| 7465c2c2-89dd-4aa6-b9af-9a61f5d38216 | VNM2 | ACTIVE | -          | Running     | net1=12.0.0.5, 172.28.11.235; sriov=11.0.0.7 |
| 2e5b990b-cbd1-4e6e-8d9b-d60c38ce435f | VNM3 | ACTIVE | -          | Running     | net1=12.0.0.6, 172.28.11.236; sriov=11.0.0.8 |
| de88aa65-49c5-4080-9468-86d3c83f8a03 | VNM4 | ACTIVE | -          | Running     | net1=12.0.0.7, 172.28.11.237; sriov=11.0.0.9 |
+--------------------------------------+------+--------+------------+-------------+-----------------------------------------------+
```

6 Testing

This chapter provides a sample testing procedure.

To verify SR-IOV configuration with OpenStack:

1. Create two VMs with different ports using the procedures in “Creating the SR-IOV Port” on page 13 and “Creating a VM with the SR-IOV Network” on page 14.

   This example creates VMs named VNM3 and VNM4.

2. Log in to the VMs using the floating IP address, which has been assigned to the private network net1 port. For the 45000 Series Adapters, ensure that the correct version of the `qede`/`qed` driver has been loaded (see “Enabling SR-IOV on the 3400/8400/45000 Series Adapters (Compute Node)” on page 6).

3. Ideally, the SR-IOV port should get DHCP IP address 11.0.0.*x* in the VM, but in this case it is not getting the IP address automatically.

   Assign an IP address to the SR-IOV port in the VMs using the `ifconfig` command, as follows:

   - ❑ VNM3—SR-IOV port IP address 11.0.0.8
   - ❑ VNM4—SR-IOV port IP address 11.0.0.9

   Following is a sample output for VNM3.

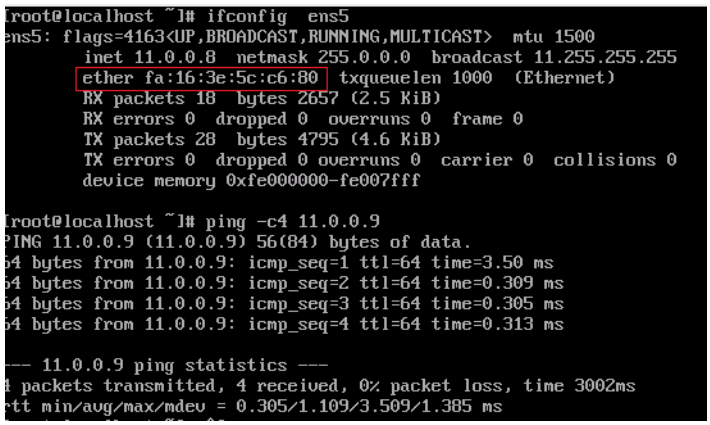

Following is a sample output for VNM4.

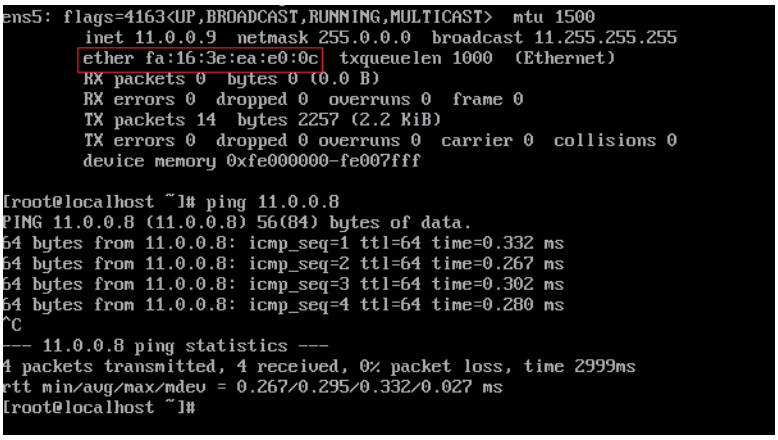

7 Known Issues

This chapter describes a known issue with OVS and OpenStack.

Sometimes OVS version 2.1.2-2 generates a segfault with the OpenStack (Kilo) release.

QLogic has installed latest version OVS-2.3.1 from the following location: http://openvswitch.org/releases/openvswitch-2.3.1.tar.gz

Following is an example of how to compile OVS.

To compile OVS on Red Hat 7.9:

1. Copy the distribution tar ball (`Openvswitch-2.3.1.tar.gz`) to the rpm source directory (`/root/rpmbuild/SOURCES`).
2. Install the following build prerequisites before compiling Open vSwitch:
   - ❑ `gcc`
   - ❑ `make`
   - ❑ `python-devel`
   - ❑ `openssl-devel`
   - ❑ `kernel-devel`
   - ❑ `graphviz`
   - ❑ `kernel-debug-devel`
   - ❑ `autoconf`
   - ❑ `automake`
   - ❑ `rpm-build`
   - ❑ `redhat-rpm-config`
   - ❑ `libtool`
3. Extract the spec file from `Openvswitch-2.3.1.tar.gz` (`Openvswitch-2.3.1/rhel/openvswitch.spec`) to the `/root/rpmbuild/SPECS/` folder.
4. Edit the `openvswitch.spec` file and remove the `Openvswitch-kmod` line from the Requires section.

5. Issue the following command:

   rpmbuild -bb /root/rpmbuild/SPECS/openvswitch.spec

   This command creates an Open vSwitch 64-bit rpm in the `/root/rpmbuild/RPMS/x86_64` location.

6. Issue the following command to install rpm:

   rpm -ivh openvswitch-2.3.1-1.x86_64.rpm

Corporate Headquarters QLogic Corporation 26650 Aliso Viejo Parkway Aliso Viejo, CA 92656 949.389.6000 **www.qlogic.com**

International Offices UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan | Israel