



# Extreme Fabric Automation OpenStack Integration Guide

Version 2.5.1

9037131-01 Rev AA  
August 2021



Copyright © 2021 Extreme Networks, Inc. All rights reserved.

## Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

## Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: [www.extremenetworks.com/company/legal/trademarks](http://www.extremenetworks.com/company/legal/trademarks)

## Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



# Table of Contents

---

<b>Preface.....</b>	<b>6</b>
Text Conventions.....	6
Documentation and Training.....	7
Help and Support.....	8
Subscribe to Product Announcements.....	8
Send Feedback.....	8
<b>About this Document.....</b>	<b>10</b>
What's New in this Document.....	10
<b>Introduction to OpenStack Integration .....</b>	<b>11</b>
Extreme Fabric Automation Overview.....	11
EFA Microservices.....	13
Fabric Service.....	13
Tenant Service.....	14
Inventory Service.....	14
Notification Service.....	14
RASlog Service.....	14
Security Service.....	14
SNMP Service.....	15
Ecosystem Integration Services.....	15
OpenStack Integration Overview.....	16
OpenStack Core Components.....	16
OpenStack Network Nodes.....	17
OpenStack Ecosystem Integration Management.....	18
OpenStack Service Commands.....	19
Limitations.....	19
Dependencies.....	20
<b>OpenStack Plugin Installation.....</b>	<b>21</b>
EFA OpenStack Plugin Package.....	21
System Requirements.....	21
Install EFA OpenStack Neutron Plugin.....	21
Upgrade or Downgrade OpenStack Neutron Plugin.....	23
Uninstall EFA OpenStack Neutron Plugin.....	23
<b>OpenStack Configuration.....</b>	<b>24</b>
Setup Overview.....	24
Topology Setup in Neutron.....	25
Tenant Creation.....	27
Configure Neutron to Connect to EFA.....	28
Enable the Mechanism Driver and Plugins.....	29
Enable the Non-Journal Plug-in.....	30

Neutron Link Mapping.....	31
ML2 Mechanism Driver.....	32
VLAN Provisioning Based on Provider Networks.....	34
Create a Single Segment Virtio VM Network.....	34
Create a Single Segment VFPT VM Network.....	37
Create a Single Segment PFPT VM Network.....	40
Create a Multi Segment Virtio VM, PFPT, and VFPT Network.....	42
L3 Service Plugin Architecture.....	44
OpenStack Network Workflow Using L3 Service Plug-in.....	45
Distributed and Centralized Routing.....	48
L3 Flavors.....	48
L3 IP MTU.....	49
L3 IPV6 ND/RA.....	50
Using Extra Routes in Neutron.....	51
Journaling.....	54
Sync EFA with Neutron.....	55
Verify EFA Health.....	57
Syslog Integration.....	58
Multiple VIM/VPOD Instances.....	58
Scale-in and Scale-out of Compute Nodes.....	61
Virtual Machine Migration.....	62
Enable VMotion.....	62
Migrate Virtual Machines.....	62
Add Certificate to OpenStack Controller.....	64
<b>EFA OpenStack Service Command Reference.....</b>	<b>65</b>
efa openstack debug.....	66
efa openstack debug router-route delete.....	67
efa openstack execution.....	68
efa openstack network show.....	69
efa openstack network-interface show.....	70
efa openstack router show.....	71
efa openstack router-interface show.....	72
efa openstack router-route show.....	72
efa openstack subnet show.....	73
efa openstack sync start.....	74
efa-blpair-mapping.....	74
efa-health show.....	76
efa-journal clear.....	78
efa-journal list.....	79
efa-journal list deps.....	79
efa-journal reset.....	81
efa-sync execute.....	82
openstack network efa-bl-pair-map create.....	84
openstack network efa-topology-link-map create.....	85
openstack network efa-topology-link-map delete.....	86
openstack network efa-topology-link-map list.....	87
openstack router create.....	88
<b>Neutron REST Endpoints.....</b>	<b>89</b>

Neutron REST Endpoints.....	89
EFA Topology Neutron Extension.....	89
EFA Border-Leaf Pair Mapping.....	90
<b>SR-IOV and Multi-Segment Support.....</b>	<b>92</b>
SR-IOV Network.....	92
Create PCI Passthrough Allowed List.....	93
Configure SR-IOV Agent.....	93
Configure Nova Scheduler.....	94
Configure Mechanism Drivers for SR-IOV.....	94
Create Network for VF-PT.....	94
Create Network for PF-PT.....	94
Create Virtual Machines.....	95
Create SR-IOV Direct Ports.....	95
Create SR-IOV Direct-Physical Port.....	95
Delete SR-IOV Entities.....	95
Multi-Segment Network.....	95
Configure Segments in Neutron.....	96
Configure Multi-Segment Network.....	96
Create Multi-Segment Network.....	97
Rename Multi-Segment Network.....	97
Create Subnet on Segment.....	97
Create a Port on SR-IOV Segment.....	97
Create a VM Using the Port.....	97



# Preface

---

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.






## Text Conventions

---

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as ExtremeSwitching switches or SLX routers, the product is referred to as *the switch* or *the router*.

**Table 1: Notes and warnings**

Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

**Table 2: Text**

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
<b>Key</b> names	Key names are written in boldface, for example <b>Ctrl</b> or <b>Esc</b> . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press <b>Ctrl+Alt+Del</b>
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
<b>NEW!</b>	New information. In a PDF, this is searchable text.

**Table 3: Command syntax**

Convention	Description
<b>bold</b> text	Bold text indicates command names, keywords, and command options.
<i>italic</i> text	Italic text indicates variable content.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ <b>x</b>   <b>y</b>   <b>z</b> }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
<b>x</b>   <b>y</b>	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member</i> [ <i>member</i> . . .].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

## Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit [www.extremenetworks.com/education/](http://www.extremenetworks.com/education/).

---

## Help and Support

---

If you require assistance, contact Extreme Networks using one of the following methods:

### Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

### The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

### Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: [www.extremenetworks.com/support/contact](http://www.extremenetworks.com/support/contact)

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

## Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

---

## Send Feedback

---

The Information Development team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.



- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, do either of the following:

- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at [documentation@extremenetworks.com](mailto:documentation@extremenetworks.com).

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.



# About this Document

---

[What's New in this Document](#) on page 10

## What's New in this Document

---

The following table describes information added to this guide for the Extreme Fabric Automation 2.5.1 software release.

**Table 4: Summary of changes**

Feature	Description	Described in
Syslog	You can now enable syslog capabilities.	<a href="#">Syslog Integration</a> on page 58
CLI commands	<ul style="list-style-type: none"><li>The new <code>efa-journal list deps</code> command dumps the journal dependency table in user-readable format.</li><li>The <code>efa-journal reset</code> command now resets the retry count for failed journal entries by specifying journal IDs.</li></ul>	<ul style="list-style-type: none"><li><a href="#">efa-journal list deps</a> on page 79</li><li><a href="#">efa-journal reset</a> on page 81</li></ul>



# Introduction to OpenStack Integration

---

[Extreme Fabric Automation Overview](#) on page 11

[EFA Microservices](#) on page 13

[OpenStack Integration Overview](#) on page 16

[OpenStack Ecosystem Integration Management](#) on page 18

[Limitations](#) on page 19

This document describes how to configure integration with the OpenStack system.

## Extreme Fabric Automation Overview

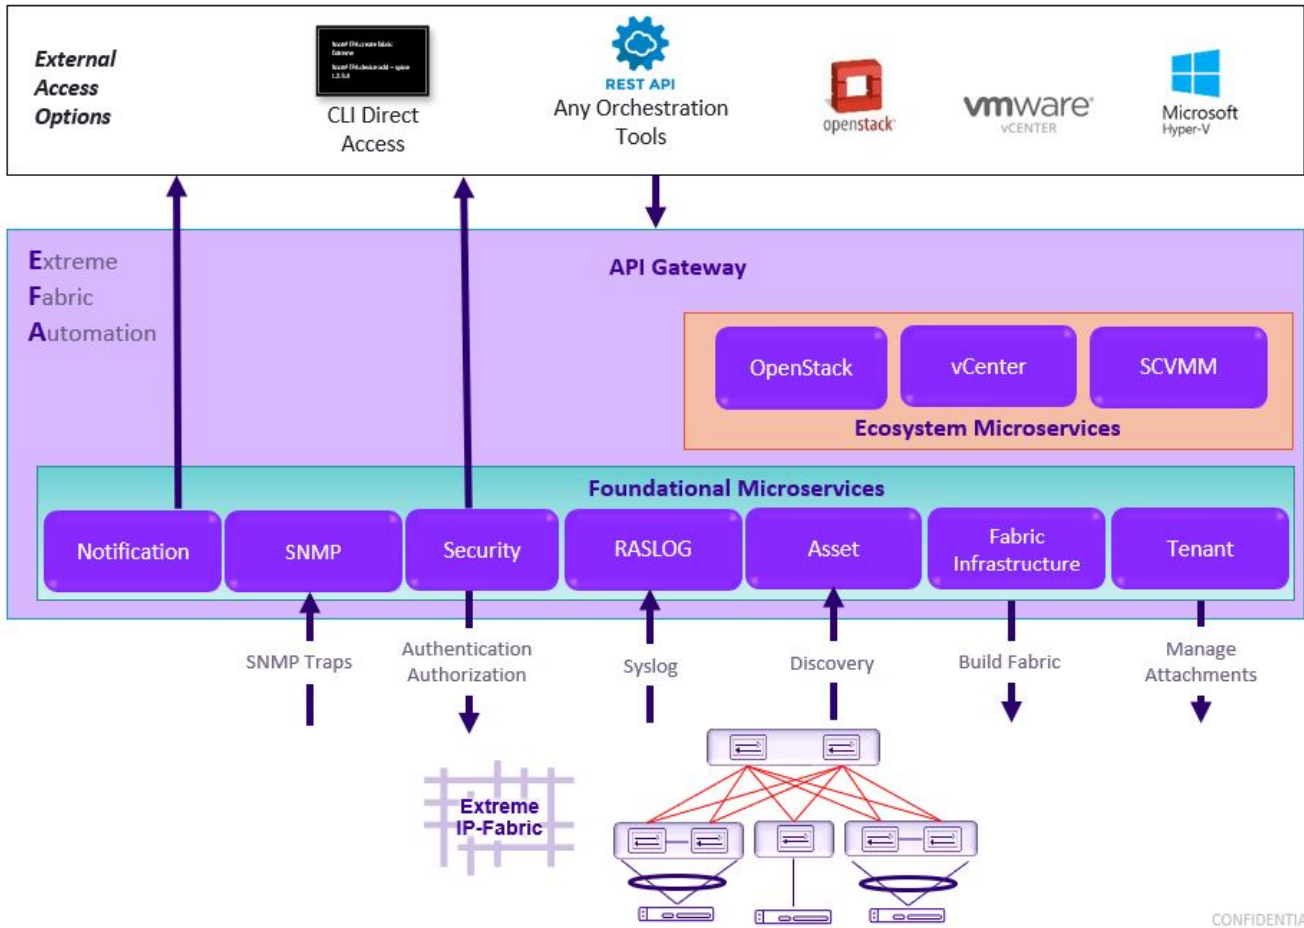
---

Extreme Fabric Automation (EFA) is a micro-services-based scalable fabric automation application.

EFA automates and orchestrates SLX IP fabrics and tenant networks, with support for the following:

- Building and managing non-Clos small data center fabrics and 3-stage and 5-stage IP Clos fabrics
- Managing tenant-aware Layer 2 and Layer 3 networks
- Configuring integration with several ecosystems: VMware vCenter, OpenStack, and Microsoft Hyper-V
- Providing a single point of configuration for your entire fabric

EFA consists of core K3s containerized services that interact with each other and with other infrastructure services to provide the core functions of fabric and tenant network automation. For more information, see [EFA Microservices](#) on page 13.




CONFIDENTIAL

Figure 1: EFA orchestration

## EFA Microservices

EFA consists of core K3s containerized microservices that interact with each other and with other infrastructure services to provide the core functions of fabric and tenant network automation.



**Figure 2: Microservices in the EFA architecture**

### Fabric Service

The Fabric Service is responsible for automating the fabric BGP underlay and EVPN overlay. By default, the EVPN overlay is enabled but you can turn it off before provisioning, if necessary. The Fabric Service exposes the CLI and REST API for automating the fabric underlay and overlay configuration.

The Fabric Service features include:

- Support for small data centers (non-Clos)
- Support for 3-stage and 5-stage Clos fabrics
- Support for MCT configuration

Underlay automation includes interface configurations (IP numbered), BGP underlay for spine and leaf, BFD, and MCT configurations. Overlay automation includes EVPN and overlay gateway configuration.

## Tenant Service

The Tenant Service manages tenants, tenant networks, and endpoints, fully leveraging the knowledge of assets and the underlying fabric. You can use the CLI and REST API for tenant network configuration on Clos and small data center fabrics.

Tenant network configuration includes VLAN, BD, VE, EVPN, VTEP, VRF, and router BGP configuration on fabric devices to provide Layer 2 extension, Layer 3 extension across the fabric, Layer 2 hand-off, and Layer 3 hand-off at the edge of the fabric.

## Inventory Service

The Inventory Service acts as an inventory of all the necessary physical and logical assets of the fabric devices. All other EFA services rely on asset data for their configuration automation. The Inventory Service is a REST layer on top of device inventory details, with the capability to filter data based on certain fields. The Inventory Service securely stores the credentials of devices in encrypted form and makes those credentials available to different components such as the Fabric and Tenant services.

The Inventory Service supports the `execute-cli` option for pushing configuration and exec commands to devices. Examples include configuring SNMP parameters or OSPF configurations. This means you can use EFA for SLX-OS commands and push the same configuration to multiple devices.

The Asset Service provides the secure credential store and deep discovery of physical and logical assets of the managed devices. The service publishes the Asset refresh and change events to other services.

## Notification Service

The Notification Service sends events, alerts, and tasks to external entities. Notifications sent from EFA are derived from the syslog events received from the devices that EFA manages. Alerts are notifications that services in EFA send for unexpected conditions. Tasks are user-driven operations or timer-based tasks such as device registration or fabric creation.

## RASlog Service

The RASlog Service processes syslog messages from devices and forwards notifications to subscribers. For more information, see RASlog Service in the [Extreme Fabric Automation Administration Guide, 2.5.0](#).

## Security Service

The Security Service consists of authentication and authorization features that enforce a security boundary between northbound clients and downstream operations between EFA and SLX devices. The service also validates users and their credentials through Role-based Access Control (RBAC) and supports local and remote (LDAP) login.



### Note

If you configure LDAP server over SSL, and use IP to connect to the server, ensure that the certificate includes the IP as part of SANs for a successful connection.

## SNMP Service

The SNMP Service processes SNMP traps from devices and forwards notifications to subscribers. For more information, see EFA as SNMP Proxy in the [Extreme Fabric Automation Administration Guide, 2.5.0](#).

## Ecosystem Integration Services

EFA provides one-touch integration with these ecosystems, providing deep insight into VMs, vSwitches, port groups, and hosts, and the translation of these into IP fabric networking constructs.

### VMware vCenter Service

The vCenter integration provides connectivity between EFA and vCenter using a REST API. EFA does not connect to individual ESXi servers. All integration is done through vCenter. For more information, see the [Extreme Fabric Automation VMware vCenter Integration Guide, 2.5.0](#). Integration support includes the following:

- Registration or deregistration of one or more vCenter servers in EFA
- Updates for vCenter asset details
- Lists of information about vCenter servers
- Inventory integration
- Dynamic updates about Tenant Service integration from vCenter and from EFA services

### Hyper-V

The Hyper-V integration supports networking configuration for Hyper-V servers in a datacenter, manual and automated configuration updates when VMs move, and visibility into the VMs and networking resources that are deployed in the Hyper-V setup. For more information, see [Extreme Fabric Automation Hyper-V Integration Guide, 2.5.0](#). Integration support includes the following:

- SCVMM (System Center Virtual Machine Manager) server discovery
- SCVMM server update
- Periodic polling of registered SCVMM servers
- SCVMM server list
- SCVMM server delete and deregister
- Network event handling

### OpenStack Service

The OpenStack service integrates Extreme OpenStack plugins with the rest of the EFA foundation services in an IP fabric. For more information, see the [Extreme Fabric Automation OpenStack Integration Guide, 2.5.0](#). Integration support includes the following:

- Create, read, update, delete (CRUD) operations on networks and ports
- LAG support
- Provider network (default, PT)
- VLAN trunking
- Network operations using single-root I/O virtualization (SR-IOV), physical and virtual functions
- vMotion (virtual machine migration)
- ML2 driver with support for:
  - Network and segment provisioning for non-default physnets

- DC-owner-based L2 extension for DC gateway
- Topology changes for port-based extension of DC gateway addition and deletion of topology entries and its changes on EFA endpoint groups.
- Single-homed connections to the edge port
- Multi-segment support
- Journaling support for L2 and L3
- L3 service plugin:
  - Routing feature support using VRF
  - Flavor (service provider) support
  - Centralized routing
  - IPv6 support (dual stack)
- Layer 3 flavors
- Neighbor Discovery and Router Advertisement support:
  - IPv6 ND MTU support
  - IPv6 No-Autoconfig support

## OpenStack Integration Overview

OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a Data Center. OpenStack Integration enables System administrators to manage and provision resources through APIs and web interface.

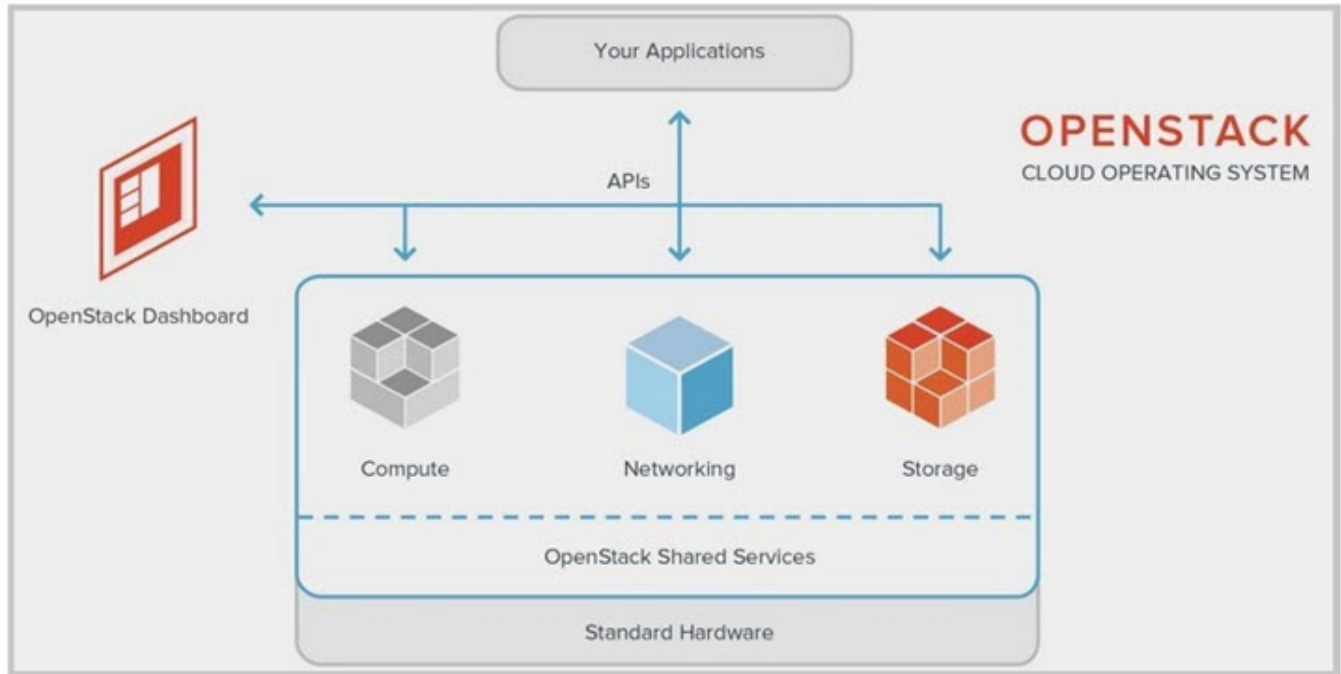
## OpenStack Core Components

The following table lists the OpenStack core components.

Component	Name	Description
Compute	Nova	Compute service that enables provisioning of compute instances or virtual servers and supports creating virtual machines and external Linux servers.
Networking	Neutron	Networking service that provides layer 2 network connectivity for virtual devices.
Dashboard	Horizon	OpenStack project that provides an extensible, unified, and web-based user interface for all OpenStack services.
Storage	Cinder	OpenStack Block Storage service for providing volumes to Nova virtual machines.

**Figure 3: OpenStack core components**





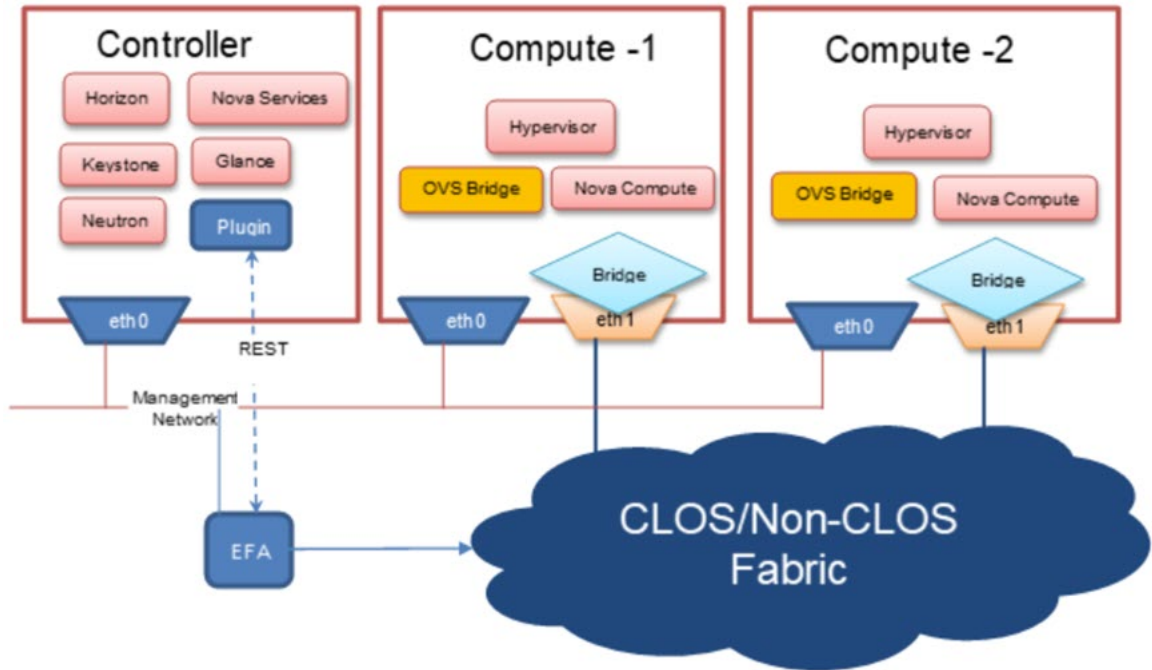
## OpenStack Network Nodes

The EFA OpenStack network consists of Controller and Compute nodes.

Most of the shared OpenStack services and other tools run on the Controller node. The Controller node supplies API, scheduling, and other shared services for the cloud. The Controller node includes dashboard, image store, and identity service. Nova Compute management service and Neutron server are also configured on the Controller node.

The VM instances or Nova Compute instances are installed on the Compute node.

The following figure shows an overview of the EFA OpenStack integration.

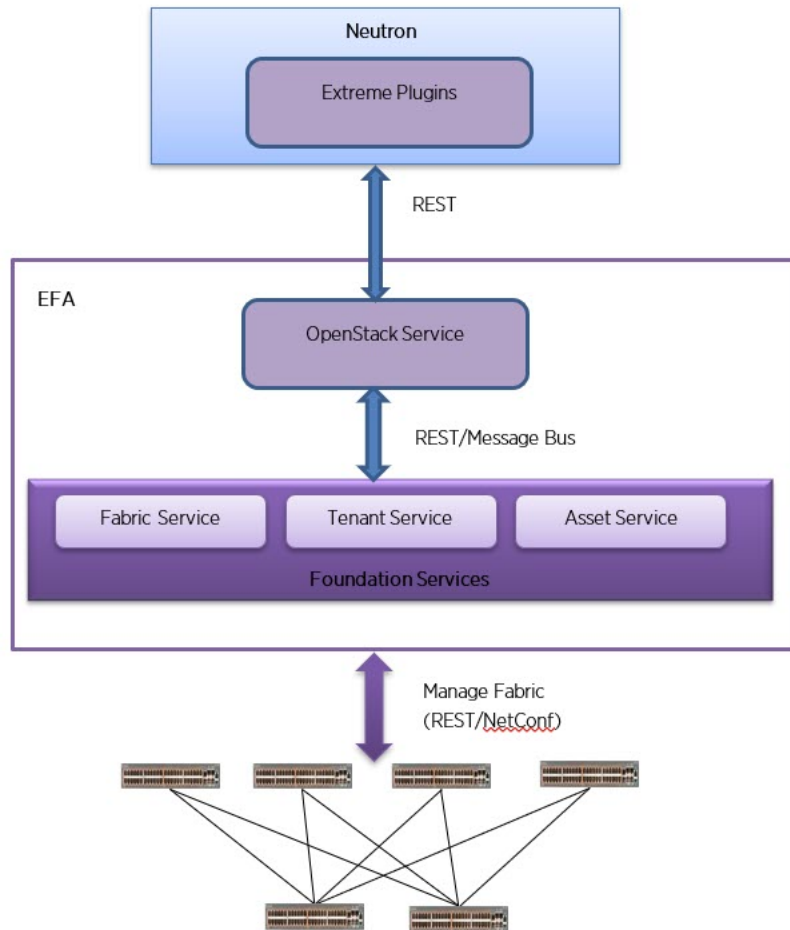


**Figure 4: Overview of EFA OpenStack Network**

## OpenStack Ecosystem Integration Management

OpenStack Service is an ecosystem service that provides integration of Extreme OpenStack plugins with the rest of the EFA foundation services in an IP fabric network.

Foundation Service	Description
Fabric Service	This service provides mechanisms to create: <ul style="list-style-type: none"> <li>• Non-Clos and Clos fabrics</li> <li>• Clos fabrics can be 3-stage or 5-stage</li> </ul>
Tenant Service	This service provides mechanisms to create: <ul style="list-style-type: none"> <li>• Layer 2 networks</li> <li>• Layer 3 networks</li> </ul>
Asset or Inventory Service	This service provides mechanisms to manage: <ul style="list-style-type: none"> <li>• Physical and logical assets of the devices in the fabric</li> <li>• The credentials of the devices</li> </ul>



**Figure 5: Overview of OpenStack ecosystem integration management**

## OpenStack Service Commands

OpenStack Service commands show the Neutron constructs and their provisioning status on EFA. For more information about the commands, see [EFA OpenStack Service Command Reference](#) on page 65.

For EFA commands and supported parameters, see the *Extreme Fabric Automation Command Reference*.

## Limitations

OpenStack ecosystem integration limitations are as follows:

- Only green field deployment is supported.
- There is an operation in Neutron to turn off routers. This operation is not supported by Extreme OpenStack support (L3 plugin or flavor), because SLX devices do not support the disabling of routers.
- EFA backup impacts OpenStack operations.

The EFA backup process is scheduled for midnight daily, by default. Any OpenStack operations that are initiated while the EFA backup is in progress can result in failures.

To recover, clean up any failed entries and retry OpenStack operations after the EFA backup process is complete.

EFA backup can be turned off or rescheduled to times that do not conflict with operations.

For more information, see the `efa system settings` command in the [Extreme Fabric Automation Command Reference, 2.5.0](#).

- Tenant VRF entities stuck in the `vrf-device-sr-delete-pending` state must be deleted for recovery. Issue the **`efa openstack debug router delete --router-id <router-id>`** command to clear the erroneous entries in the tenant. The **`efa-sync execute`** command on the Neutron side synchronizes the routers, router-ports, and route entries to the EFA.

## Dependencies

OpenStack ecosystem integration dependencies are as follows:

- Non-Clos (2-node, 4-node, and 8-node)
- Clos
  - 3-stage Clos (leaf or spine nodes)
  - 5-stage Clos (leaf, spine, or super-spine nodes)
- Compute nodes (connected to every leaf node in MCT and non-MCT configurations)



# OpenStack Plugin Installation

---

[EFA OpenStack Plugin Package](#) on page 21

[Install EFA OpenStack Neutron Plugin](#) on page 21

[Upgrade or Downgrade OpenStack Neutron Plugin](#) on page 23

[Uninstall EFA OpenStack Neutron Plugin](#) on page 23

This chapter provides steps for installing the OpenStack EFA plugin.

## EFA OpenStack Plugin Package

---

OpenStack integration requires the Modular Layer 2 (ML2) Mechanism Driver to interact with Extreme Fabric Automation (EFA). The Neutron drivers and plugins communicate with EFA over the REST interface.

The EFA plugin for OpenStack is packaged as RPM files.

The EFA OpenStack plugin package supports many features. For details, see [OpenStack Ecosystem Integration Management](#) on page 18.

## System Requirements

System requirements for EFA are provided in the *Extreme Fabric Automation Deployment Guide*.

## Install EFA OpenStack Neutron Plugin

---

This section provides information required to install Extreme Fabric Automation OpenStack Neutron plugin on Ubuntu.

### Before You Begin

The prerequisites for installing the EFA OpenStack Neutron plugin are as follows:

- Working knowledge of Linux
- Experience in OpenStack deployment
- Experience in managing EFA 2.x.x
- Network nodes are prepared as required
- All OpenStack compute, network, and controller node hostname name resolutions are setup (fully qualified Host names (fqdn) are not supported for beta release)
- All OpenStack nodes are configured with unique hostnames
- Working EFA Fabric using `efa cli/rest`

- OpenStack nodes are connected to the leaf switches either in direct mode, VPC, or bonded mode
- Bonding setup is done using 802.3ad



### Note

EFA OpenStack Neutron plugin supports only Extreme specific OpenStack services. If other OpenStack services are required, install the respective plugins prior to EFA OpenStack Neutron plugin installation.

All network and server connection settings can be saved to .csv files for configuring them simultaneously using the file option.

## Procedure

1. Install the EFA OpenStack plugin RPM package.

```
# $sudo rpm -U <RPM file>
```

2. Note the Neutron configuration file layout.

- Neutron configuration: /etc/neutron/neutron.conf
- ML2 plugin configuration: /etc/neutron/plugins/ml2/ml2\_conf.ini
- Extreme EFA Mechanism driver or topology configuration: /etc/neutron/plugins/ml2/ml2\_conf\_extreme.ini

3. Configure the M12 core\_plugin in the neutron.conf file.

```
[DEFAULT]
core_plugin=m12
service_plugins = trunk, segments, efa_topology_plugin
```

Do not enable the reference router plugin.

4. Enable Neutron in the ml2\_conf\_extreme.ini file to communicate with EFA.

```
[ml2_extreme]
efa_rest_token = <efa_api_token for VIM_1>
efa_cert_file = /root/gcla/extreme-ca-chain.crt
efa_secure_mode = True
efa_port = 443
efa_host = efa.extremenetworks.com
region_name = VIM_1
region_shared = SHARED_TENANT
#SHARED_TENANT is the name of the shared tenant created on EFA
fabric_name = CNCF
```

5. Enable the Neutron EFA extension plugin in the ml2\_conf\_extreme.ini file to build initial physical topology between OpenStack Compute nodes and TOR switches.

```
[efa_topology]
efa_link_mapping_file = /home/ubuntu/link.csv
```

6. Enable Extreme EFA mechanism drivers in ml2\_conf.ini.

```
ml2_conf.ini
[ml2]
tenant_network_types = vlan
type_drivers = vlan
mechanism_drivers = openvswitch,extreme_efa
[ml2_type_vlan]
network_vlan_ranges = physnet1:100:500 (Required vlan range)
[ovs]
bridge_mappings = physnet1:br0 (bridge used for datapath)
```

7. Modify the system unit file to start Neutron with `m12_conf_extreme.ini`.

```
# ExecStart = /usr/local/bin/neutron-server --config-file /etc/neutron/neutron.conf --  
config-file /etc/neutron/plugins/m12/m12_conf.ini --config-file /etc/neutron/  
plugins/m12/m12_conf_extreme.ini
```

```
# systemctl daemon-reload
```

On DevStack installation, modify the `/etc/systemd/system/devstack@q-svc.service` file.

8. Restart the Neutron server.

```
systemctl restart devstack@q-agt.service
```

along with

```
devstack@q-svc.service
```

If you are installing using Open Source, use the `sudo service neutron-* restart` command.

If you are installing using DevStack, use the `sudo systemctl restart devstack@q-svc.service` command.

9. Verify if the status of the Neutron server is `Active` and confirm the following:

- The Neutron service was started with the `m12_conf_extreme.ini` file.
- The **efa-topology** extension is loaded using `openstack extension show efa-topology`.

```
# sudo systemctl status devstack@q-svc.service
```

If you are installing using Open Source stack, use the `sudo service neutron-* status` command.

## Upgrade or Downgrade OpenStack Neutron Plugin

### Procedure

Upgrade or downgrade the EFA plugin RPM package from the network node.

```
# $sudo rpm -U <RPM file>
```

## Uninstall EFA OpenStack Neutron Plugin

### Procedure

Uninstall the EFA OpenStack Neutron plugin.

```
# sudo rpm -e networking-extreme
```



# OpenStack Configuration

---

- [Setup Overview](#) on page 24
- [ML2 Mechanism Driver](#) on page 32
- [L3 Service Plugin Architecture](#) on page 44
- [Journaling](#) on page 54
- [Sync EFA with Neutron](#) on page 55
- [Verify EFA Health](#) on page 57
- [Syslog Integration](#) on page 58
- [Multiple VIM/VPOD Instances](#) on page 58
- [Scale-in and Scale-out of Compute Nodes](#) on page 61
- [Virtual Machine Migration](#) on page 62
- [Add Certificate to OpenStack Controller](#) on page 64

This chapter describes how to configure OpenStack for use with EFA.

## Setup Overview

---

The following summarizes the phases in configuring OpenStack to work with EFA. Details are provided in the sections that follow.

The instructions assume that the IP Fabric has already been set up using Fabric Service commands. For details on creating a Clos or non-Clos IP Fabric, refer to the EFA Administration Guide.

**Table 5: Phases in Configuring OpenStack for EFA**

Steps	Explanation
Create fabric	The instructions in the following sections assume that the IP Fabric has already been set up.
Set up physical topology	Shows SLX switch topology and its interconnection with compute nodes
Create tenant	Setting up the EFA tenants according to the physical topology
Configure Neutron	Configuration of Neutron to connect to EFA
Enable mechanism driver and plugins	Enabling Layer2, Layer3 plugins and drivers on Neutron



**Table 5: Phases in Configuring OpenStack for EFA (continued)**

Steps	Explanation
Enable journaling	Alternate Neutron configuration for setting up Layer2 and Layer3 using journaling
Map Neutron links	Link mapping according to the physical topology on Neutron

## Topology Setup in Neutron

Use OpenStack CLI to manage the topology between compute nodes and SLX devices. The Physical Network Topology extension provides a Neutron CLI for the OpenStack administrator to manage links between SLX devices and Compute NICs. Neutron has an abstract for physical network called Provider Network.

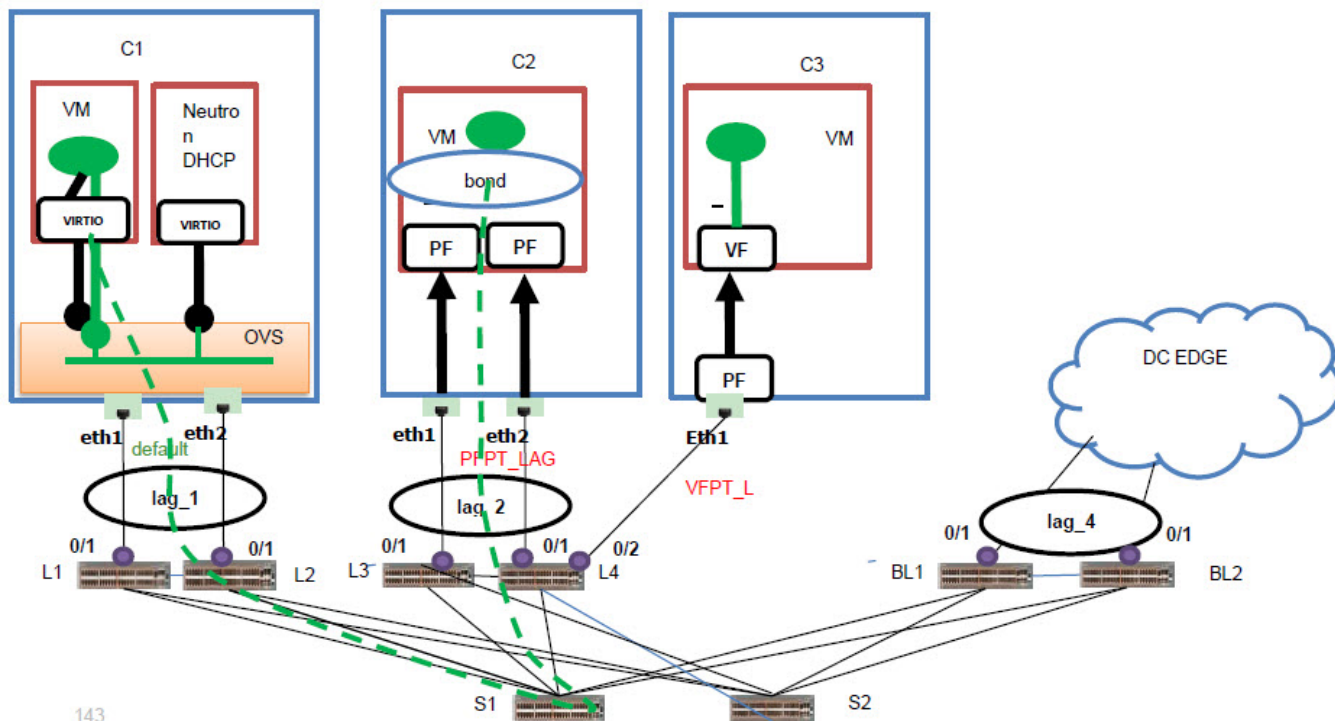
The Extreme Topology plugin provides **efa-link-mapping** command to configure and set up the topology. Use the efa-link-mapping feature to scale in to and scale out of compute nodes. For more information about scaling in and scaling out of compute nodes, see [Scale-in and Scale-out of Compute Nodes](#) on page 61.



### Note

All network and server connection settings and mappings can be saved to `.csv` files for configuring them simultaneously using the `startup` file option in the `m12_conf_extreme.ini`.

Network provisioning and de-provisioning on EFA depends on physnet mapping for non-default physnets. When a node is added or deleted from the mapping file, the EFA operations are updated accordingly.



**Figure 6: Topology setup in Neutron**

Use the settings in the follow table to set up the physical topology.

Node	IP Address or Description
C1	10.24.51.115 (OpenStack115)
C2	10.24.51.116 (OpenStack116)
C3	10.24.51.117 (OpenStack117)
L1	10.24.14.133
L2	10.24.14.134
L3	10.24.14.135
L4	10.24.14.136
BL1	10.24.14.191
BL2	10.24.14.192

## Tenant Creation

One OpenStack instance is mapped to an EFA tenant. This section describes the steps involved in setting up the tenant on EFA.

**Table 6: Steps for setting up tenant on EFA**

Steps	Description
Create a Tenant	Create an EFA tenant. The name of the tenant should be the same as the <b>region_name</b> in Neutron.
Create Lags/Port Channel	Create lag according to the physical topology. The names of the lags must be used in defining the topology mapping on Neutron.
Create Shared Tenant	Create an EFA tenant that contains the shared resources across multiple EFA tenants. The name of the tenant must be used as <b>region_shared</b> in ML2 configurations.
EFA REST TOKEN	Used as API key <b>efa_rest_token</b> in ML2 configurations for secure connection towards EFA.

The following are the EFA CLI commands used to achieve the steps above:

**Table 7: EFA CLI commands for setting up tenant on EFA**

EFA CLI Command	
<code>efa tenant create --name <b>VIM_1</b> --vlan-range 2-4090 --l2-vni-range 1-5000 --port 10.24.14.133.[0/1], 10.24.14.134.[0/1], 10.24.14.135.[0/1], 10.24.14.136.[0/1-2]</code>	Create Tenant VIM_1
<code>efa tenant po create --name lag_1 --tenant <b>VIM_1</b> --port 10.24.14.133.[0/1], 10.24.14.134.[0/1] --speed 10Gbps --negotiation active</code>	Lag on C1
<code>efa tenant po create --name lag_2 --tenant <b>VIM_1</b> --port 10.24.14.135.[0/1], 10.24.14.136.[0/1] --speed 10Gbps --negotiation active</code>	Lag on C2
<code>efa tenant create --name <b>SHARED_TENANT</b> --port 10.24.14.191.[0/1], 10.24.14.192.[0/1] -type <b>shared</b></code>	Shared tenant on border leaf
<code>efa tenant po create --name lag_4 --tenant <b>SHARED_TENANT</b> --port 10.24.14.133.[0/1], 10.24.14.134.[0/1] --speed 10Gbps --negotiation active</code>	Border leaf lag created on shared tenant

**Table 7: EFA CLI commands for setting up tenant on EFA (continued)**

EFA CLI Command	
<code>efa auth client register --name VIM_1 --type openstack</code>	Register <b>VIM_1</b> tenant for OpenStack access
<code>efa auth apikey generate --client-id &lt;client_id&gt;</code>	Use the Client ID generated in client register for API key generation Use the generated API Key as a the <b>efa_rest_token</b> in OpenStack

## Configure Neutron to Connect to EFA

The entities in the following script need to be configured for the OpenStack instance to connect to EFA.

### Procedure

Edit `/etc/neutron/plugins/ml2/ml2_conf_extreme.ini` on VIM-1:

```
[ml2_extreme]
efa_rest_token = <efa_api_token for VIM_1>
efa_cert_file = /root/gcla/extreme-ca-chain.crt
efa_secure_mode = True
efa_port = 443
efa_host = efa.extremenetworks.com
use_fqdn = false
region_name = VIM_1
region_shared = SHARED_TENANT
#SHARED_TENANT is the name of the shared tenant created on EFA
fabric_name = CNCF
[efa_topology]
efa_link_mapping_file = /home/ubuntu/link.csv
```

The EFA tenant name and the `region_name` in OpenStack should be the same.

### Token Description

Token	Description
<code>efa_rest_token</code>	API token available using <code>efa client</code> commands
<code>efa_cert_file</code>	The location of the CA Cert Chain for EFA
<code>efa_secure_mode</code>	True (for secure connections)
<code>efa_port</code>	Port running the HTTP service
<code>efa_host</code>	IP address of EFA service: <code>efa.extremenetworks.com</code> . This setting must be resolved to the EFA IP address through DNS or <code>/etc/hosts</code> .
<code>region_name</code>	Used as the tenant name by EFA
<code>region_shared</code>	Name of the shared tenant as created on EFA
<code>efa_link_mapping_file</code>	Mapping of physical NICs to Extreme switch ports

## Enable the Mechanism Driver and Plugins

Make configurations on the Neutron controller for starting the Extreme EFA Mechanism Driver and L3 Service plugin, a journal-based plugin.

### About This Task

The `extreme_efa_v2` and `extreme_l3_efa_v2` plugins use a journaling mechanism,

### Procedure

1. **ML2 Configuration:** Set the mechanism driver to `extreme_efa_v2`, along with `openvswitch`. In the `/etc/neutron/plugins/ml2/ml2_conf.ini` file, add the following lines:

```
[ml2]
mechanism_drivers = openvswitch,extreme_efa_v2,sriovnicswitch
```

The configurations in the following table also affect the working of the journaling mechanism, which is configured in `/etc/neutron/plugins/ml2/ml2_conf.ini`.

Parameter	Description	Default
<code>sync_timeout</code>	Sync thread timeout in seconds or fraction.	10
<code>retry_count</code>	Number of times to retry a row before failing.	5
<code>maintenance_interval</code>	Journal maintenance operations interval in seconds.	300
<code>completed_rows_retention</code>	Time to keep completed rows (in seconds). For performance reasons, do not change this from the default value (0) which indicates completed rows are not kept. This value is checked every <code>maintenance_interval</code> by the cleanup thread. To keep completed rows indefinitely, set the value to -1.	0
<code>journal_rows_threshold</code>	Maximum number of journal rows allowed. If the number of rows exceeds this threshold, then the <code>efa_health</code> tool shows the status as down.	1000
<code>efa_status_interval</code>	Journal EFA reachability check interval in seconds	150

2. **Topology Plugin Configuration:** Activate the topology plugin by adding the following lines in the `/etc/neutron/neutron.conf` file:

```
[DEFAULT]
service_plugins = efa_topology_plugin,trunk,segments
```

3. **L3 Service Plugin Configuration:** Activate the L3 Service plugin by adding the following lines in the `/etc/neutron/neutron.conf` file:

```
[DEFAULT]
service_plugins = extreme_l3_efa_v2,efa_topology_plugin,trunk,segments
```



#### Note

`extreme_l3_efa_v2` is an L3 Service plugin and is a replacement for the default L3 Service plugin (called `router`). Only one of them should be used.

4. **L3 Flavor Configuration:** Enable L3 flavors with Extreme by adding the relevant service providers to `neutron.conf`:

```
[service_providers]
service_provider = L3_ROUTER_NAT:extremev2:networking_extreme.l3.l3_flavor_v2
.ExtremeL3ServiceProviderv2:default
```

Flavors require the default L3 Service plugin (called `router`). In other words, flavors can be used only when the service plugin is `router`.

```
service_plugins = router,efa_topology_plugin,trunk,segments
```

For more information, see [L3 Flavors](#) on page 48.

## Enable the Non-Journal Plug-in

This is an alternative Neutron configuration for setting up Layer 2 and Layer 3 without using journaling.

Complete the following changes in the configuration files to enable the non-journal plug-in.

### *ML2 Configuration*

The following configuration lines must be available in `/etc/neutron/plugins/ml2/ml2_conf.ini`. The mechanism driver must be set to `extreme_efa` along with `openvswitch`.

```
[ml2]
mechanism_drivers = openvswitch,extreme_efa,sriovnicswitch
```

### *L3 Service Plug-in Configuration*

The following configuration lines must be available in the `/etc/neutron/neutron.conf` file for the L3 Service plug-in to be activated.

```
[DEFAULT]
service_plugins = extreme_l3_efa,efa_topology_plugin,trunk,segments
```



#### Note

`extreme_l3_efa` is an L3 Service Plug-in and is a replacement for the default L3 Service Plug-in (called `router`). Only one of them should be used.

### *L3 Flavor Configuration*

To enable L3 flavors with Extreme, service providers must be added to `neutron.conf`, as follows. Multiple service providers flavor, that is `ExtremeL3ServiceProvider` and `ExtremeL3ServiceProviderv2`, can be added to the configuration file, but only one can be the

default. The router creation command must specify the flavor profile for whichever router is being created. For more information, see [L3 Flavors](#) on page 48.

```
[service_providers]
service_provider =
L3_ROUTER_NAT:extreme:networking_extreme.l3.l3_flavor. ExtremeL3ServiceProvider:default
```



### Note

Flavors require the default L3 Service Plug-in (called `router`). In other words, flavors can be used only when the service plug-in is `router`:

```
service_plugins =
  router,efa_topology_plugin,trunk,segments
```

## Neutron Link Mapping

You can configure the topology entries on Neutron according to the physical topology.

The following console script can be used for configuring the topology entries on Neutron according to the physical topology. These commands provide mapping of host, NIC, provider-network to switch, port, lag. If the lag or po is not created on EFA, then the `--po-name` should be left empty.

```
efa-link-mapping add --host Openstack115 --nic eth1 --pn default --switch 10.24.14.133 --
port 0/1 --po-name lag_1
efa-link-mapping add --host Openstack115 --nic eth2 --pn default --switch 10.24.14.134 --
port 0/1 --po-name lag_1
efa-link-mapping add --host Openstack116 --nic eth1 --pn PFPT_LAG --switch 10.24.14.135 --
port 0/1 --po-name lag_2
efa-link-mapping add --host Openstack116 --nic eth2 --pn PFPT_LAG --switch 10.24.14.136 --
port 0/1 --po-name lag_2
efa-link-mapping add --host Openstack117 --nic eth1 --pn VFPT_L --switch 10.24.14.136 --
port 0/2
efa-link-mapping add --host DC-GW1 --pn EXT1 --switch 10.24.14.191 --
port 0/1 --po-name lag_4
efa-link-mapping add --host DC-GW1 --pn EXT1 --switch 10.24.14.192 --
port 0/1 --po-name lag_4
```

The same information can be provided as part of the startup file which is specified in `ml2_conf_extreme.ini`.

```
opensuse@Openstack114:~$ cat /home/opensuse/link.csv
Openstack115,eth1,default,10.24.14.133,0/1,lag_1
Openstack115,eth2,default,10.24.14.134,0/1,lag_1
Openstack116,eth1,PFPT_LAG,10.24.14.135,0/1,lag_2
Openstack116,eth2,PFPT_LAG,10.24.14.136,0/1,lag_2
Openstack117,eth1,VFPT_L10.24.14.136,0/1
```

```
DC-GW1, ,EXT1,10.24.14.191,0/1,lag_4
DC-GW1, ,EXT1,10.24.14.192,0/1,lag_4
```

**Note**

The number of comma delimiters used in each line within the CSV file must be the same. This information can be provided through a `-file` option in the command.

```
efa-link-mapping add --file /home/opensuse/link.csv
```

Fields	Description
Host	Host name of the compute node (as seen in OpenStack server list). In case of a DC Gateway (border leaf), this will be a label like DC-GW1.
nic	NIC on the compute node. In case of a DC Gateway (border leaf), this is empty.
pn	Neutron physnet that is being mapped
switch	IP address of the physical switch
Port	Switch interface that is connected to the NIC. In case of a DC Gateway (border leaf), this will be the switch interface that is connected to an external DC Edge.
po-name	Port channel name that was already created using the tenant command.

Only VLAN-based Physnet entries should be added to this table.

**Note**

The same behavior can be achieved using the `openstack network efa-topology-link-map` command.

### *Topology Mapping CLI Commands*

The following commands are available for showing and manipulating topology mapping:

- `openstack network efa-topology-link-map list`
- `openstack network efa-topology-link-map create`
- `openstack network efa-topology-link-map delete`

For more information, see [openstack network efa-topology-link-map list](#) on page 87 and the related commands.

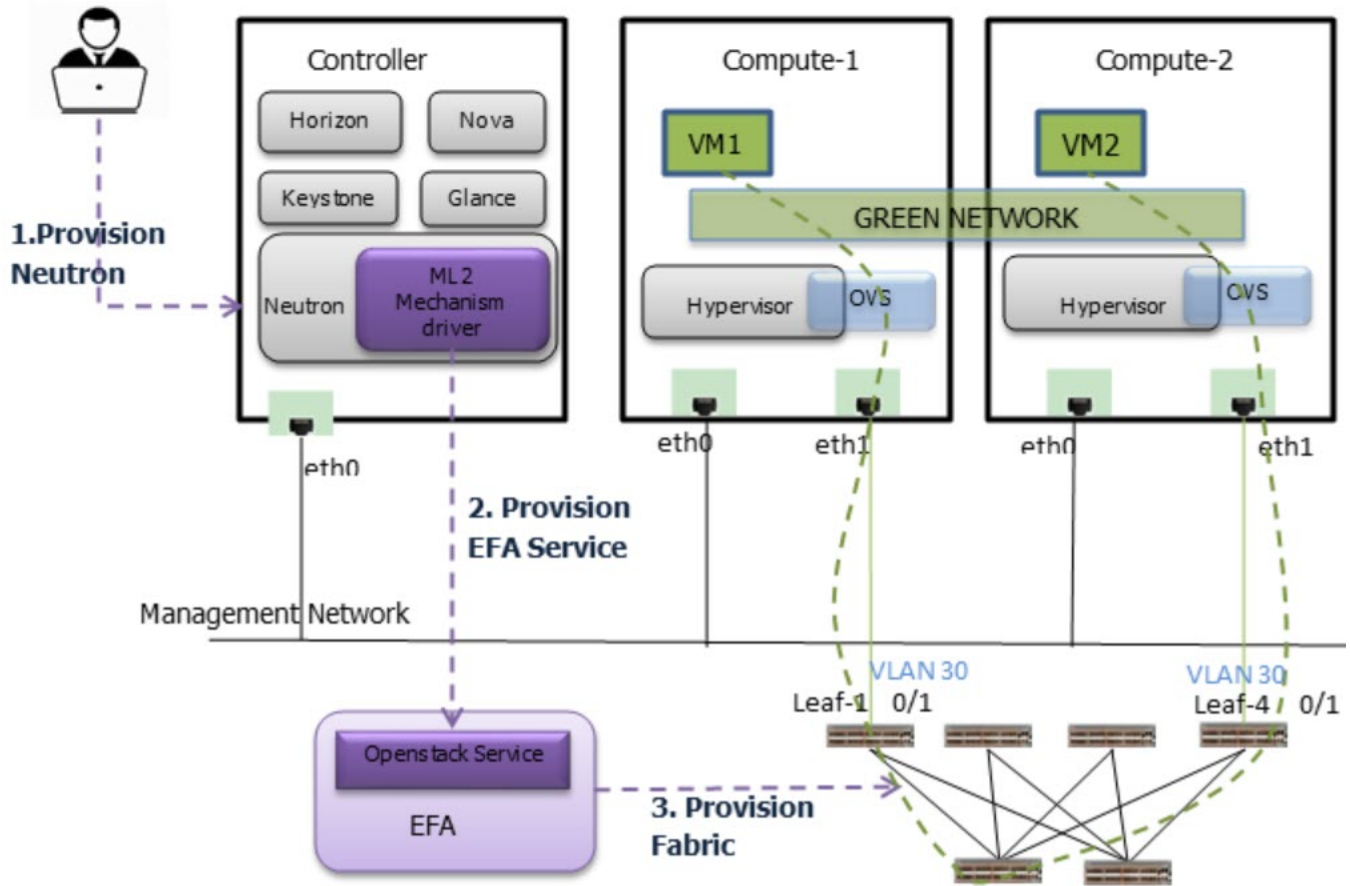
## ML2 Mechanism Driver

The Modular Layer 2 (ml2) plugin is a framework allowing OpenStack networking to simultaneously utilize the variety of layer 2 networking technologies found in complex real-world data centers.



External networks are managed using ml2 Mechanism Driver. The mechanism driver is responsible for taking the information established by the type driver and ensuring that it is properly applied, given the specific networking mechanisms that are enabled.

The following figure shows an overview of the ML2 mechanism driver within Neutron.



**Figure 7: Overview of ML2 Mechanism Driver within Neutron**

Extreme ML2 Mechanism driver within the Neutron component of OpenStack proxies the Neutron API calls for network management toward the OpenStack Service running in EFA.

The OpenStack Service in EFA translates the neutron network management calls to appropriate tenant API calls and provision the fabric with appropriate L2 networking constructs.

## VLAN Provisioning Based on Provider Networks

There are several scenarios in which Virtual Machine can be created on provider networks.

**Table 8: VLAN provisioning based on provider networks (physnet)**

Provider Network (Physnet)	VLAN Provisioning by ML2
default	VLAN provisioning of endpoints is done when Neutron ports are bound to a host or compute node. Example: <ul style="list-style-type: none"> <li>• The Neutron port is bound to a host when VM is launched.</li> <li>• The Neutron port is bound to a host (Neutron controller) when DHCP starts.</li> </ul>
default (device-owner=dc-edge -host=<DCGW>)	VLAN provisioning for default physnet is done during port creation based on additional parameters passed during port create call.
non-default (PFPT_L, EXT1)	VLAN provisioning is done during network creation or segment creation of single segment or multiple segment.

The endpoints are deprovisioned during the following negation operations:

- `virtio` Ports are unbounded from a host
- `port delete` operations on `device-owner=dc-edge` qualified ports
- Deletion of networks or segments on non-default physnets



### Note

The ML2 driver supports overlapping VLAN across provider physical networks.

Example

```
openstack network create --provider-network-type vlan --provider-physical-
network physnet1
--provider-segment 3530 phynetwork1
openstack network create --provider-network-type vlan --provider-physical-
network default
--provider-segment 530 defNetwork
```

## Create a Single Segment Virtio VM Network

You can create a single segment Virtio VM network on a default physnet and extend the network to DC Edge.




### Note

VLAN provisioning depends on the following:

- Neutron Virtio ports are bound to a host during VM launch.
- DHCP ports are bound during subnet creation.
- DC Edge extension is achieved using explicit `port-create`.

The following figure shows an overview of a single segment Virtio VM Network with DC Edge.



**Figure 8: Overview of single segment Virtio VM network**

**Table 9: Commands and impacts**

Command	EFA Impact
<pre>openstack network create -- provider-network-type vlan -- provider-physical-network default --provider-segment 3320 sslnetwork1</pre>	<p>EPG Created for sslnetwork Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d CTAG = 3320</p> <p><b>Note:</b> 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d is the neutron UUID allocated for the EPG.</p>
<pre>openstack subnet create sslsubnet1 --network sslnetwork1 --subnet- range 10.1.1.0/24 openstack subnet create sslsubnet1ip6 --network sslnetwork1 --ip-version 6 --ip6- address-mode=dhcpv6-stateful -- subnet-range fd00:10:0:57::1000/64</pre>	<p>DHCP EndPoints Created on EPG corresponding to sslnetwork1. <b>VLAN Provisioned.</b> EPG Updated Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d Port = lag_1 (added)</p>

**Table 9: Commands and impacts (continued)**

Command	EFA Impact
<pre>openstack port create sslVirtIoTrunkPort1 --network sslnetwork1 openstack network trunk create -- parent-port sslVirtIoTrunkPort1 sslVirtIoTrunk1 openstack port create sslVirtIoSubport1 --network sslnetwork1</pre>	
<pre>openstack network create -- provider-network-type vlan -- provider-physical-network <b>default</b> --provider-segment <b>3321</b> sslnetwork2</pre>	<p>EPG Created for sslnetwork2  Name = 84cbf489-f3d9-41c7-bbb2-6cb7df33da6d  CTAG = 3321</p>
<pre>openstack subnet create sslsubnet2 --network sslnetwork2 --subnet- range 11.1.1.0/24 openstack subnet create sslsubnet2ipv6 --network sslnetwork2 --ip-version 6 --ipv6- address-mode=dhcpv6-stateful -- subnet-range fd00:11:0:57::1000/64</pre>	<p>DHCP EndPoints Created on EPG corresponding to sslnetwork2. <b>VLAN Provisioned</b>  EPG Updated  Name = 84cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Port = lag_1 (added)</p>
<pre>openstack port create sslVirtIoTrunkPort2 --network sslnetwork2 openstack network trunk create -- parent-port sslVirtIoTrunkPort2 sslVirtIoTrunk2 openstack network trunk set -- subport port=sslVirtIoSubport1,segmentation -type=vlan,segmentation-id=3801 sslVirtIoTrunk2</pre>	
<pre>openstack server create --flavor m1.large --image ubuntu --port \$ (neutron port-list   grep -w 'sslVirtIoTrunkPort1'   awk '{print \$2}') sslVirtIoVM1 --availability- zone nova:Openstack116</pre>	<p>Endpoint corresponding to 'sslVirtIoTrunkPort1'  added to EPG(sslnetwork1) <b>VLAN Provisioned</b>  EPG Updated  Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Port = lag_1 (added) - no impact already added</p>
<pre>openstack server create --flavor m1.large --image ubuntu --port \$ (neutron port-list   grep -w 'sslVirtIoTrunkPort2'   awk '{print \$2}') sslVirtIoVM2 --availability- zone nova:Openstack117</pre>	<p>Endpoint corresponding to 'sslVirtIoTrunkPort2'  added to EPG(sslnetwork2) <b>VLAN Provisioned</b>  EPG Updated  Name = 84cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Port = lag_1 (added) - no impact already added  Endpoint corresponding to sslVirtIoSubport1  added to EPG(sslnetwork1) <b>VLAN Provisioned</b>  EPG Updated  Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Port = lag_1 (added) - no impact already added</p>

**Table 9: Commands and impacts (continued)**

Command	EFA Impact
<pre>openstack port create ss2DcGwPort1 --device-owner network:dc_edge -- host <b>DCGW-1</b> --network ss1network1</pre>	EndPoint corresponding to 'host DCGW-1' added to EPG (ss1network1) <b>VLAN Provisioned</b> EPG Updated <i>Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d</i> <i>Port = lag_1, lag_4 (added)</i>
<pre>openstack port create ss2DcGwPort1 --device-owner network:dc_edge -- host <b>DCGW-1</b> --network ss1network2</pre>	EndPoint corresponding to 'host DCGW-1' added to EPG (ss1network2) <b>VLAN Provisioned</b> EPG Updated <i>Name = 84cbf489-f3d9-41c7-bbb2-6cb7df33da6d</i> <i>Port = lag_1, lag_4 (added)</i>

## Create a Single Segment VFPT VM Network

You can create a single segment VFPT VM network on a non-default physnet and extend the network to DC Edge.

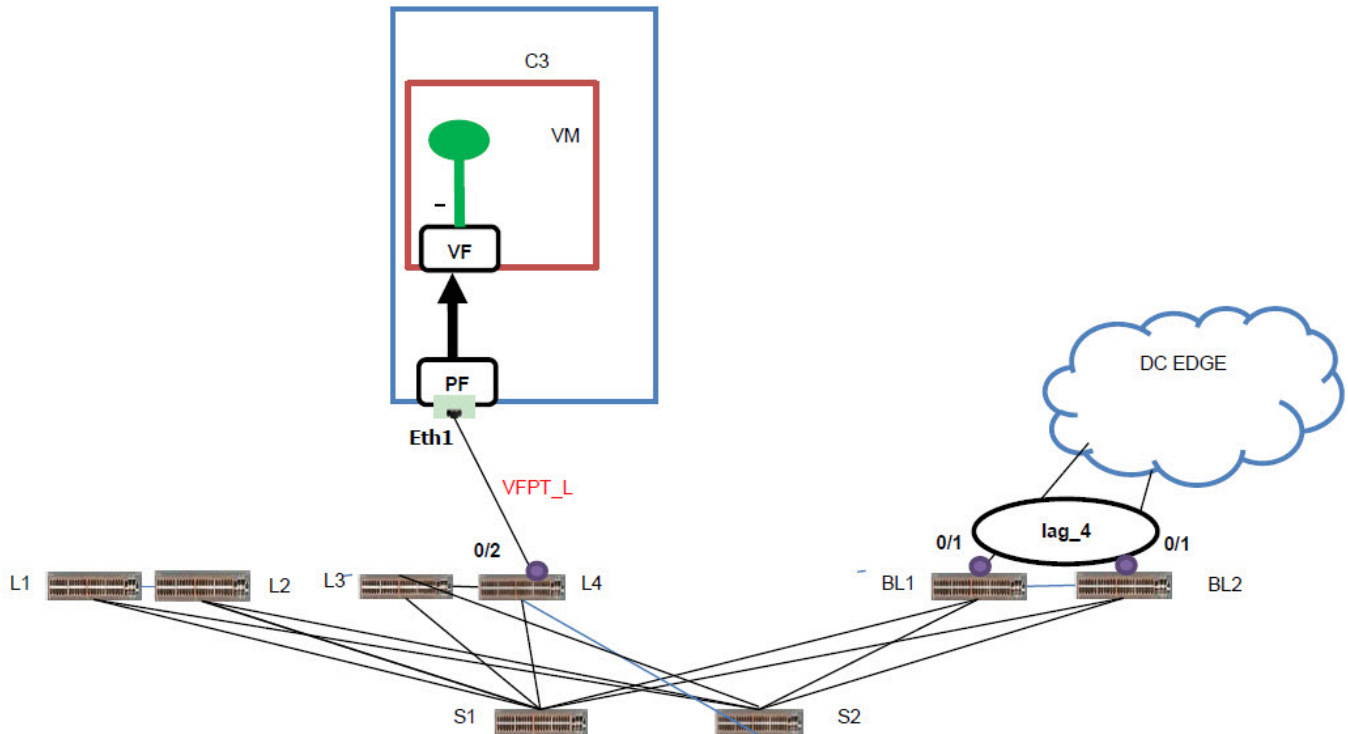


### Note

VLAN provisioning depends on the following:

- VFPT ports are bound to a host when the VFPT subport is attached to the trunk port.
- DC Edge extension is achieved using explicit port-create.

The following figure shows an overview of a single segment VFPT VM network with DC Edge. The non-default physnet in this scenario is VFPT\_L.



**Figure 9: Overview of single segment VFPT VM network with DC Edge**

**Table 10: Commands and impacts**

Command	Impact on EFA
<pre>openstack network create -- provider-physical-network VFPT_L -- provider-network-type flat ss9_vfpt_flat_left</pre>	No impact, as operations are on a FLAT network type.
<pre>openstack network create -- provider-physical-network VFPT_R -- provider-network-type flat ss9_vfpt_flat_right</pre>	
<pre>openstack subnet create ss8flatleftsubnet --network ss9_vfpt_flat_left --no-dhcp -- subnet-range 107.1.1.0/24</pre>	
<pre>openstack port create --network ss9_vfpt_flat_left --vnic-type direct ss9_port_vfpt_left1 openstack network trunk create -- parent-port ss9_port_vfpt_left1 ss9SriovTrunkLag1</pre>	


**Table 10: Commands and impacts (continued)**

Command	Impact on EFA
<pre>openstack server create --flavor myhuge --image ubuntu --port \$ (neutron port-list   grep -w 'ss9_port_vfpt_left1'   awk '{print \$2}') ss9SriovVM1 --availability- zone nova:compute-0-10.domain.tld --poll</pre>	
<pre>openstack network create -- provider-network-type vlan -- provider-physical-network VFPT_L -- provider-segment <b>3390ss9network</b></pre>	<p>EPG Created for <b>ss9network</b>  Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  CTAG = <b>3390</b>  <b>Note</b> - 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Is the neutron UUID allocated for the EPG</p>
<pre>openstack subnet create ss9subnet1 --network ss9network1 --no-dhcp -- subnet-range 90.1.1.0/24 openstack subnet create ss9subnetlipv6 --network ss8network1 --ip-version 6 --no- dhcp --subnet-range fd00:90:0:57::1000/64</pre>	No impact, as the --no-dhcp option is used.
<pre>openstack port create ss9SriovSubPort1 --network ss9network1 --mac-address &lt;same- mac-as-ss9_port_vfpt_left1&gt; --vnic- type direct --fixed-ip subnet=ss9subnet1,ip- address=90.1.1.10 --fixed-ip subnet=ss9subnetlipv6,ip- address=fd00:90:0:57::10</pre>	
<pre>openstack network trunk set -- subport port=ss9SriovSubPort1,segmentation- type=vlan,segmentation-id=3390 ss9SriovTrunkLagopenstack port set ss9SriovSubPort1 --device-owner compute:nova --host compute-0-10.domain.tld --device &lt;same-deviceid-as- ss9_port_vfpt_left1&gt; --binding- profile pci_slot=&lt;same-as-slot-of- ss9_port_vfpt_left1&gt; --binding- profile pci_vendor_info=&lt;same-as- vendor-of-ss9_port_vfpt_left1&gt; -- binding-profile physical_network=&lt;same-as-physnet- of-ss9_port_vfpt_left1&gt;</pre>	<p>Endpoint corresponding to ss9SriovSubPort1  added to EPG(ss9network1) <b>VLAN Provisioned</b>  EPG Updated  Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Port = L4[0/2] (added)</p>
<pre>openstack port create ss9DcGwPort --network ss9network1 --device- owner network:dc_edge --host <b>DCGW-1</b></pre>	<p>EndPoint corresponding to 'host DCGW-1' added  to EPG (ss9network1) <b>VLAN Provisioned</b>  EPG Updated  Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Port = L4[0/2]. lag_4 (added)</p>

## Create a Single Segment PFPT VM Network

You can create a single segment PFPT VM network on a non-default physnet and extend the network to DC Edge.

The following figure shows an overview of a single segment PFPT VM network with DC Edge. The “non-default” physnet in this scenario is PFPT\_LAG.



**Figure 10: Overview of single segment PFPT VM network with DC Edge**



**Table 11: Commands and impacts**

Command	EFA Impact
<pre> openstack network create -- provider-physical-network PFPT_L -- provider-network-type flat ss7_pfpt_flat_left openstack network create -- provider-physical-network PFPT_R -- provider-network-type flat ss7_pfpt_flat_right openstack subnet create ss7flatleftsubnet --network ss7_pfpt_flat_left --no-dhcp -- subnet-range 107.1.1.0/24 openstack subnet create ss7flatrightsubnet --network ss7_pfpt_flat_right --no-dhcp -- subnet-range 108.1.1.0/24 openstack port create --network pfpt_flat_left --vnic-type direct- physical ss7_port_pfpt_left openstack port create --network pfpt_flat_right --vnic-type direct- physical ss7_port_pfpt_right openstack network trunk create -- parent-port ss7_port_pfpt_left ss7PFPTTrunkLag1 openstack server create --flavor myhuge --image ubuntu --port \$( neutron port-list   grep -w 'ss7_port_pfpt_left'   awk '{print \$2}') --port \$(neutron port-list   grep -w 'ss7_port_pfpt_right'   awk '{print \$2}') ss7PFPTLAGVM1 -- availability-zone nova:compute-0-10.domain.tld --poll </pre>	No impact, as operations are on a FLAT network type.
<pre> openstack network create -- provider-network-type vlan -- provider-physical-network PFPT_LAG --provider-segment <b>3370ss7network1</b> </pre>	<p>EPG Created for <b>ss7network1</b>  Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  CTAG = <b>3370</b>  <b>Note</b> - 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d  Is the neutron UUID allocated for the EPG</p>
<pre> openstack subnet create ss7subnet1 --network ss7network1 --no-dhcp -- subnet-range 70.1.1.0/24 openstack subnet create ss7subnetlipv6 --network ss7network1 --ip-version 6 --no- dhcp --subnet-range fd00:70:0:57::1000/64 </pre>	No impact as -no-dhcp option is used.

**Table 11: Commands and impacts (continued)**

Command	EFA Impact
<pre>openstack port create ss7PFPTSubPort1 --network ss7network1 --mac-address &lt;same- mac-as-ss7_port_pfpt_left&gt; --vnic- type direct-physical --fixed-ip subnet=ss8subnet1, ip- address=70.1.1.10 --fixed-ip subnet=ss7subnet1ipv6, ip- address=fd00:70:0:57::10</pre>	
<pre>openstack network trunk set -- subport port=ss7PFPTSubPort1, segmentation- type=vlan, segmentation-id=3370</pre>	<p>Endpoint corresponding to ss7PFPTSubPort1 added to EPG(ss7network1) <b>VLAN Provisioned</b>                      EPG Updated  <i>Name = 84cbf489-f3d9-41c7-bbb2-6cb7df33da6d</i>  <i>Port = lag_2 (added)</i></p>
<pre>openstack port create ss7DcGwPort --network ss7network1 --device- owner network:dc_edge --host DCGW-1 --fixed-ip subnet=ss7subnet1, ip- address=70.1.1.30 --fixed-ip subnet=ss7subnet1ipv6, ip- address=fd00:70:0:57::30</pre>	<p>EndPoint corresponding to 'host DCGW-1' added to EPG (ss9network1) <b>VLAN Provisioned</b>                      EPG Updated  <i>Name = 84cbf489-f3d9-41c7-bbb2-6cb7df33da6d</i>  <i>Port = lag_2, lag_4 (added)</i></p>

## Create a Multi Segment Virtio VM, PFPT, and VFPT Network

You can create default physnet and segments for PFPT\_LAG and VFPT\_L.

Perform this procedure to create default physnet and segments for PFPT\_LAG and VFPT\_L. There are Virtual Machine created on all of these segments. The network also gets extended on DC EDGE using segments.

The following figure shows an overview of Virtio VM, PFPT, and VFPT segment network with DC Edge.

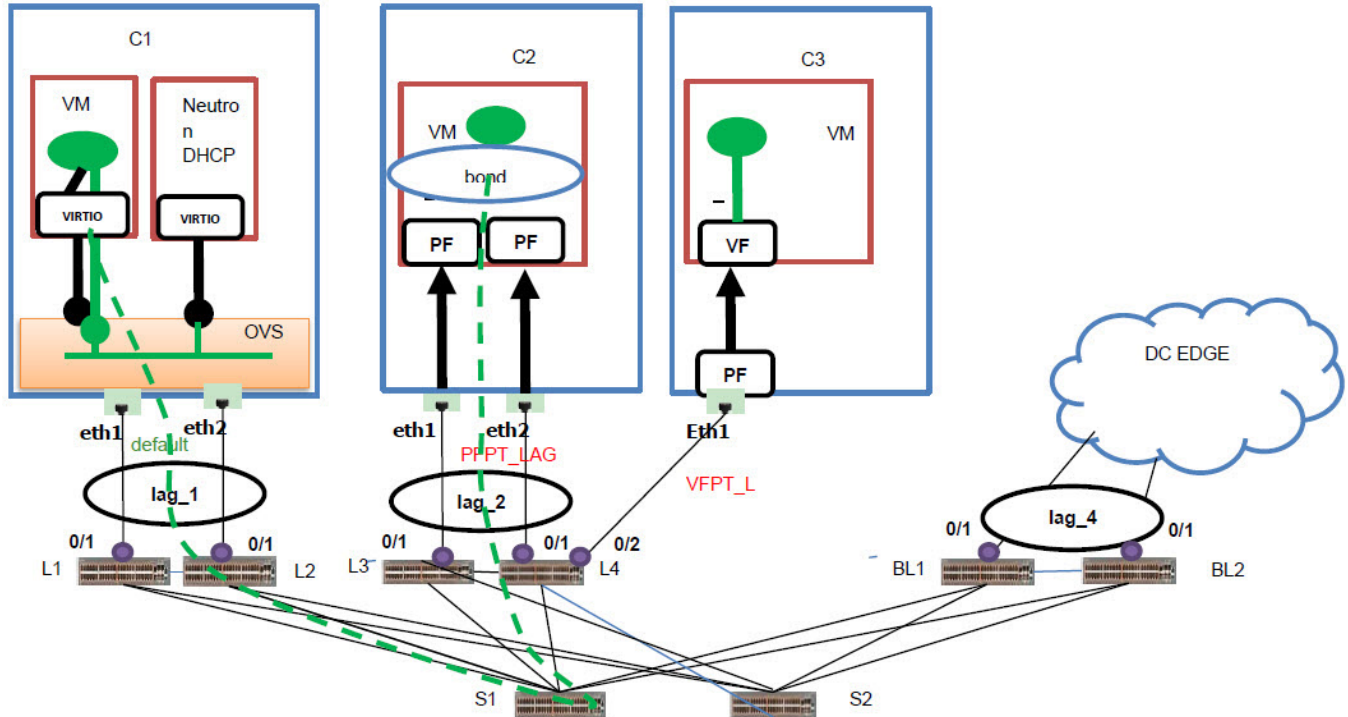


Figure 11: Overview of Virtio VM, PFPT, and VFPT segment network with DC Edge

Table 12:

Command	EFA Impact
<pre>openstack network create -- provider-network-type vlan -- provider-physical-network VFPT_L -- provider-segment 3710 ms10Network1</pre>	EPG Created for ms10Network1 Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d CTAG = 3710 <b>Note</b> - 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d is the neutron UUID allocated for the EPG Port = L4[0/2] added
<pre>openstack network segment create -- network-type vlan --physical- network EXT1 --segment 3710 -- network ms10Network1 ms10Network1DcgwSegment</pre>	Endpoint corresponding EXT1 added to EPG(ms10Network1) <b>VLAN Provisioned</b> EPG Updated Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d Port = L4[0/2], <b>lag_4</b> (added)
<pre>openstack network segment create -- network-type vlan --physical- network PFPT_LAG --segment 3710 -- network ms10Network1 ms10Network1PFPTSegment</pre>	Endpoint corresponding EXT1 added to EPG(ms10Network1) <b>VLAN Provisioned</b> EPG Updated Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d Port = L4[0/2], <b>lag_4</b> , <b>lag_2</b> (added)
<pre>openstack network segment create -- network-type vlan --physical- network default --segment 3710 -- network ms10Network1 ms10Network1DefaultSegment</pre>	

**Table 12: (continued)**

Command	EFA Impact
<pre>openstack subnet create ms10subnet1 --network ms10Network1 --subnet- range 130.1.1.0/24 openstack subnet create ms10subnetlipv6 --network ms10Network1 --ip-version 6 --ipv6- address-mode=dhcpv6-stateful -- subnet-range fd00:0130:0:57::1000/64</pre>	
<pre>openstack port create ms10VirtIoPort1 --network ms10Network1 --vnic-type normal openstack port create ms10SriovPort2 --network ms10Network1 --vnic-type direct openstack port create ms10PFPTport3 --network ms10Network1 --vnic-type direct-physical</pre>	
<pre>openstack server create --flavor myhuge --image ubuntu --port \$ (neutron port-list   grep -w 'ms10VirtIoPort1'   awk '{print \$2}') ms10VirtIoVM1 --availability- zone nova:compute-0-5.domain.tld</pre>	<p>Endpoint corresponding default added to EPG(ms10Network1) <b>VLAN Provisioned</b> EPG Updated <i>Name = 74cbf489-f3d9-41c7-bbb2-6cb7df33da6d</i> <i>Port = L4[0/2], lag_4, lag_2 (added)</i></p>
<pre>openstack server create --flavor myhuge --image ubuntu --port \$ (neutron port-list   grep -w 'ms10SriovPort2'   awk '{print \$2}') ms10SrIovVM2 --availability- zone nova:compute-0-1.domain.tld openstack server create --flavor myhuge --image ubuntu --port \$ (neutron port-list   grep -w 'ms10PFPTport3'   awk '{print \$2}') ms10PFPTVM3 --availability-zone nova:compute-0-7.domain.tld</pre>	

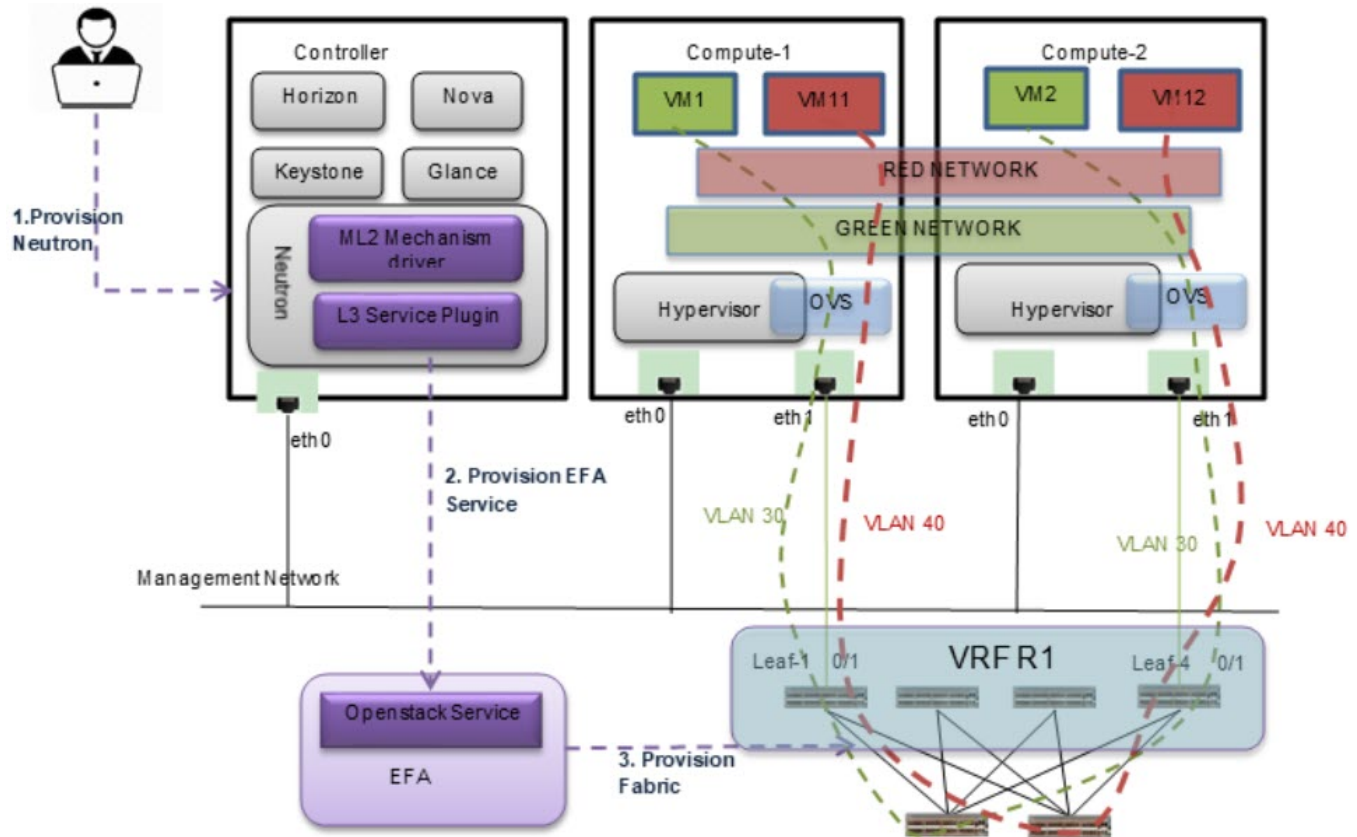
## L3 Service Plugin Architecture

The Extreme L3 Service plugin within the Neutron component of OpenStack proxies the Neutron API calls for network management toward the OpenStack Service running in EFA.

The OpenStack Service in EFA translates the Neutron network management calls to appropriate tenant API calls and provision the fabric with appropriate L3 networking constructs.

The L3 Service plugin can create a distributed or centralized VRF instance on the IP fabric which corresponds to the Neutron router settings.

For distributed routers, VRF instances are created on corresponding leaf devices. For centralized routers, VRFs are created on border leaf devices only. For more details, see [Distributed and Centralized Routing](#) on page 48.



**Figure 12: L3 Service plugin**

## OpenStack Network Workflow Using L3 Service Plug-in

### About This Task

Extreme ML2 Plugin and L3 Service Plugin proxies these requests towards EFA with VLAN details about the allocated by the type driver.

Openstack Service on EFA utilizes the End Point Group (EPG) construct of Tenant Services to provision an End Point Group (EPG).

### Procedure

1. Create an OpenStack network named GREEN\_NETWORK using a Neutron CLI command.

```
openstack network create GREEN_NETWORK (assume VLAN 30 is allocated by type driver)
openstack subnet create GREEN_SUBNET_IPV4 --subnet-range 10.0.0.0/24 --network GREEN_NETWORK
```

2. Create a Virtual Machine VM1 on the Compute-1 attached to GREEN\_NETWORK.

```
openstack server create --nic
    net-id=$(neutron net-list | awk '/GREEN_NETWORK/ {print $2}') --image
cirros-0.3.4-x86_64-uec --flavor m1.tiny
--availability-zone nova:Compute-1 VM1
```

3. Create a Virtual Machine VM2 on the Compute-2 attached to GREEN\_NETWORK.

```
openstack server create --nic
    net-id=$(neutron net-list | awk '/GREEN_NETWORK/ {print $2}') --image
cirros-0.3.4-x86_64-uec --flavor m1.tiny --availability-zone
nova:Compute-2 VM2
```

4. Create a second OpenStack network by the name RED\_NETWORK using Neutron CLI command.

```
openstack network create RED_NETWORK (assume
    VLAN 30 is allocated by type driver)
openstack subnet create RED_SUBNET_IPV4 --subnet-range 9.0.0.0/24 -network
RED_NETWORK
```

5. Create a Virtual Machine VM11 on the Compute-1 attached to RED\_NETWORK.

```
openstack server create --nic
    net-id=$(neutron net-list | awk '/ RED_NETWORK / {print $2}') --image
cirros-0.3.4-x86_64-uec --flavor m1.tiny
--availability-zone nova:Compute-1 VM11
```

6. Create a Virtual Machine VM12 on the Compute-2 attached to RED\_NETWORK.

```
openstack server create --nic
    net-id=$(neutron net-list | awk '/ RED_NETWORK / {print $2}') --image
cirros-0.3.4-x86_64-uec --flavor m1.tiny
--availability-zone nova:Compute-2 VM12
```

7. Create a Router R1 and add the two networking instances, GREEN\_SUBNET and RED\_SUBNET as part of the Router.

```
openstack router create R1
openstack router add subnet R1 GREEN_SUBNET_IPV4
openstack router add subnet R1 RED_SUBNET_IPV4
```

## Results

Extreme ML2 Plugin and L3 Service Plugin proxies these requests towards EFA with VLAN details about the allocated by the type driver.

OpenStack Service on EFA utilizes the End Point Group (EPG) construct of Tenant Services to provision an End Point Group (EPG).



### Note

You can also create IPv6 subnet using the following command:

```
openstack subnet create --ip-version 6 --network GREEN_NETWORK --subnet-range
10:2000::/64
--gateway 10:2000::1 GREEN_SUBNET_IPV6
openstack router add subnet R1 GREEN_SUBNET_IPV6
```

Port-based addition of router interface is supported.

Create extra routes for the L3 networks using the command:

```
openstack router add route --route destination= 10.4.4.0/24,gateway=10.0.0.2 R1
```



#### Note

Using the same command, you can also create IPv6 static routes.

**Table 13: EFA impact of L3 provisioning**

VM	Neutron Network	Tenant Service(EPG)
VM1	GREEN_NETWORK(VLAN30) UUD =74cbf489-f3d9-41c7- bbb2-6cb7df33da6d	74cbf489-f3d9-41c7-bbb2-6cb7df33da6ds <ul style="list-style-type: none"> <li>Endpoint: eth 0/1 on Leaf-1</li> <li>EndPoint: eth 0/1 on Leaf-2</li> </ul>
VM2	GREEN_NETWORK(VLAN30) UUD =74cbf489-f3d9-41c7- bbb2-6cb7df33da6d	<ul style="list-style-type: none"> <li>CTAG 30</li> <li><b>AnyCast</b> 10.0.0.1</li> <li><b>VRF R1(UUID=99</b>cbf489-f3d9-41c7- bbb2-6cb7df33da03)</li> </ul>
VM11	RED_NETWORK(VLAN40) GREEN_NETWORK(VLAN30) UUD =89cbf489-f3d9-41c7- bbb2-6cb7df33da02	89cbf489-f3d9-41c7-bbb2-6cb7df33da02 <ul style="list-style-type: none"> <li>Endpoint: eth 0/1 on Leaf-1</li> <li>EndPoint: eth 0/1 on Leaf-2</li> <li>CTAG 40</li> </ul>
VM12	RED_NETWORK(VLAN40) UUD =89cbf489-f3d9-41c7- bbb2-6cb7df33da02	<ul style="list-style-type: none"> <li><b>Anycast</b> 9.0.0.1</li> <li><b>VRF R1(UUID=99</b>cbf489-f3d9-41c7- bbb2-6cb7df33da03)</li> </ul>

Router R1 will be realized as VRF in EFA tenant services.

```
Name                : 99cbf489-f3d9-41c7-bbb2-6cb7df33da03
Tenant              : Region_7_8
Routing Type       : centralized
Centralized Routers : 10.20.246.8
                   : 10.20.246.7
Redistribute       : connected
Max Path           : 8
Local Asn          :
L3VNI              :
EVPN IRB BD        :
EVPN IRB VE        :
BR VNI             : 4096
BR BD              :
BR VE              :
RH Max Path        :
Enable RH ECMP     : false
Enable Graceful Restart : false
Route Target       : import 101:101
                   : export 101:101
Static Route       : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                   : 10.20.246.8->10.4.4.0/24,10.0.0.2
                   : 10.20.246.7->10.4.4.0/24,10.0.0.2
Static Route BFD   :
```

EPG provisioning on the fabric creates an L2 network on the fabric spanning VM1, VM2, VM11, and VM12 with necessary fabric mappings. This creates the necessary constructs to establish an end-to-end

connectivity between the two sets of Virtual Machines VM1 and VM2, VM11 and VM12. The VRF configuration enables routing between the two networks.

## Distributed and Centralized Routing

L3 routers can be configured in centralized mode or distributed mode.

In centralized mode, routing is configured only on the border leaf pairs. In contrast, for distributed mode, routing is configured on the corresponding leaf nodes where the endpoints reside.

During router creation, You can provide centralized mode or distributed mode as input:

```
openstack router create R1 --distributed
openstack router create R2 --centralized
```

The default option is centralized, so the following command creates centralized routing:

```
openstack router create R2
```

The L3 service plugin passes this information to EFA. Then, the EFA Tenant Service configures VRF or routing on border leafs or on leaf nodes, based on the configuration mode.



### Note

- The OpenStack integration works with only one pair of border leaf devices. The centralized routing instance is created on this pair. Ensure that you add only one border leaf pair during fabric creation.
- Only centralized router is supported at this time. The creation of a distributed router is disabled.

## L3 Flavors

L3 flavors allow multiple backends in one Neutron deployment.

### About This Task

Each backend has its own logical network topology, completely separate from all others. This architecture allows incremental migrations from one backend to another.

When you create a router with a specific flavor, the flavor framework looks up the service profiles for that flavor and creates the router using the associated drivers. Once a driver is selected for a router, the router or driver association is stored in the database, so any future operations on the router can look up the driver without going through the flavor framework.

You can use the following command to create a centralized router along with the flavor: `neutron router-create router1 --flavor l3_extreme_v2 -distributed=False`.



For changes required in the Neutron configuration file, see [Configure Neutron to Connect to EFA](#) on page 28 and [Enable the Non-Journal Plug-in](#) on page 30.



#### Note

- The **openstack router create** command does not support the creation of routers using flavors.
- The functionality to configure distributed mode is disabled. For more information, see [Distributed and Centralized Routing](#) on page 48.

### Procedure

1. Create a flavor profile <flavorprofileid> by issuing `openstack network flavor profile create --driver networking_extreme.l3.l3_flavor_v2.ExtremeL3ServiceProviderv2`.
2. Create a network flavor using `openstack network flavor create --service-type L3_ROUTER_NAT l3_extreme_v2`.
3. Add a profile to the flavor using `openstack network flavor add profile l3_extreme_v2 <flavorprofileid>`.
4. Create a Neutron router using the create flavor. Issue the command `neutron router-create router1 --flavor l3_extreme_v2`.
5. Create a flavor profile with non-journaling support using `openstack network flavor profile create --driver networking_extreme.l3.l3_flavor.ExtremeL3ServiceProvider`.

## L3 IP MTU

Neutron provides multiple ways to specify the maximum transmission unit (MTU) for networks.

You can provide a global value in the `neutron.conf` file using the `global_physnet_mtu` value. This value applies to all networks. For example:

```
/etc/neutron/neutron.conf
[DEFAULT]
global_physnet_mtu = <mtu_value>
```

You can provide a specific value when creating the network using the `--mtu` option. This value takes higher precedence and overrides the global value. For example:

```
openstack network create --mtu <mtu-value> GREEN_NETWORK
```

The MTU value for each network is captured by Extreme plug-ins and passed to EFA. EFA then enforces the value on the endpoint group as part of its `ipv6-nd-mtu` and `ip-mtu` options. This value is per Ctag and is enforced on the VE interface corresponding to the Ctag. EFA configures the `ipv6-nd-mtu` and `ip-mtu` values on the specific VE interface on the device.

The value is applied to the SLX device when the router-interface is associated with the network in Neutron, either when the router is added to a subnet that is associated with the network or when a port on the network is added to a router.

The following is a sample SLX configuration.

```
(device)# show running-config interface Ve 3

interface Ve 3
  vrf forwarding 274e9461081849f8867b77a112f6c510
  ip anycast-address 61.1.1.1/24
  ip mtu 9000
  ipv6 anycast-address fd00:61:0:57::1/64
  ipv6 nd managed-config-flag
  ipv6 nd other-config-flag
  ipv6 nd mtu 9000
  ipv6 nd prefix fd00:61:0:57::/64 2592000 604800 no-autoconfig
  no shutdown
```



#### Note

SLX supports IP MTU values in the range from 1280 through 9194 bytes. Creation of the router-interface fails if you try to configure an MTU value outside of this range.

## L3 IPV6 ND/RA

When creating subnets in Neutron, you can provide values to specify IPv6 Router Advertisement mode and IPv6 address mode.

Values to specify IPv6 Router Advertisement mode and IPv6 address mode can be one of: `dhcpv6-stateful`, `dhcpv6-stateless` or `slaac` (stateless address autoconfiguration).

The subnet settings take effect when it is added to a router.

Example:

```
openstack subnet create --ip-version 6 --subnet-range fd00:10:0:57::/64
--gateway fd00:10:0:57::1000 --network net1 subnet1 --ipv6-address-mode dhcpv6-stateful
--ipv6-ra-mode dhcpv6-stateful

openstack router subnet add Router1 subnet1
```

Based on the value passed in `ipv6-ra-mode` and `ipv6-address-mode`, the A, M, O bits are set in the IPv6 Router Advertisements sent from the external router.

- A - autoconfig flag
- M - managed-config-flag
- O - Other-config-flag

The following table explains the router advertisement flag settings and expected guest instance behavior.

**Table 14: Router advertisement flag settings and expected guest instance behavior**

S number	ipv6 ra mode	ipv6 address mode	A, M, O bits set in Router Advertisements sent from Extreme Router	Guest instance behavior
1	slaac	slaac	1, 0, 0	Guest instance obtains IPv6 address from non-OpenStack router using SLAAC.
2	dhcpv6-stateful	dhcpv6-stateful	0, 1, 1	Guest instance obtains IPv6 address from dnsmasq using DHCPv6 stateful and optional information from dnsmasq using DHCPv6.
3	dhcpv6-stateless	dhcpv6-stateless	1, 0, 1	Guest instance obtains IPv6 address from non-OpenStack router using SLAAC and optional information from dnsmasq using DHCPv6.



#### Note

Other combinations for `ipv6-ra-mode` and `ipv6-address-mode` in the above command are considered invalid.

Sample SLX configuration:

```
Leaf4# show running-config interface ve
interface Ve 1330
vrf forwarding vrf_2052
ipv6 anycast-address fd00:10:0:57::1000/64
ipv6 nd managed-config-flag
ipv6 nd other-config-flag
ipv6 nd prefix fd00:10:0:57::/64 2592000 604800 no-autoconfig
no shutdown
!
```

- When the M bit is set to 1, `managed-config-flag` is set on the VE interface
- When the O bit is set to 1, `other-config-flag` is set on the VE interface.
- When the A bit is set to 0, `no-autoconfig` is added to the `ipv6 nd prefix` command on the VE interface.

## Using Extra Routes in Neutron

Extreme Networks supports static routes using extra routes in Neutron.

The Extreme Neutron L3 Router plugin and flavors implement the extra routes API that enables the addition of static routes via OpenStack. An extra route is added using the `destination` and `gateway` fields, representing the destination CIDR and the next hop gateway IP address. Deletion of

routes can also be performed with equivalent REST API and CLI commands. Routes installed on the router can be listed with the `openstack router show` command.

The support is available for both IPv4 and IPv6 address families.

Desc	Description
<code>openstack router add route --route destination=20.0.4.0/24,gateway=10.35.40.1 myRouter</code>	Creates route in router myRouter with destination CIDR and gateway IP.
<code>openstack router remove route --route destination=20.0.4.0/24,gateway=10.35.40.1 myRouter</code>	Removes route in router myRouter with matching destination CIDR and gateway IP.
<code>openstack router show myRouter</code>	Shows the router properties, including list of routes installed in the router
<code>openstack router add remove route --route destination=2002:35:40::0/64,gateway=2001:35:40::1 myRouter</code>	IPv6 address family destination CIDR and gateway IP can be provided



#### Note

- You can add a duplicate route, that is, one with the same destination CIDR and gateway, or delete a non-existing route to or from a router.
- There is no field available in the `openstack router show` command to display route status. Even if the static routes fail to configure on the fabric side, their status is not shown in OpenStack.
- Extra routes are supported only in `extreme_v2` drivers with journaling support.

#### Border leaf pair selection

An EFA command for border leaf pair mapping lets you specify the border leaf pair on which the route will be installed.

#### Creating a border leaf pair

```
openstack network efa-bl-pair-map create --name <string> --switch-ips <comma separated string of ips>
```

#### Example

```
openstack network efa-bl-pair-map create --name blp1 --switch-ips 10.2.2.2,10.2.2.3
+-----+-----+
| Field      | Value                                |
+-----+-----+
| id         | 0d19abc5-d31b-4682-b86d-cc1617ea0c97 |
| name       | blp1                                  |
| switch_ips | 10.2.2.2,10.2.2.3                    |
+-----+-----+
```

You can also add the border leaf pair by using a topology link mapping file. In the topology link mapping border leaf pair, mappings can be provided in the new section `[efa-bl-pairs]`. This

section should be provided after topology link mapping entries. If more than one entry is provided, then only the first entry will be considered.

#### Example

```
opensuse@Openstack114:~$ cat /home/opensuse/link.csv
Openstack115,eth1,default,10.24.14.133,0/1,lag_1
Openstack115,eth2,default,10.24.14.134,0/1,lag_1
Openstack116,eth1,PFPT_LAG,10.24.14.135,0/1,lag_2
Openstack116,eth2,PFPT_LAG,10.24.14.136,0/1,lag_2
Openstack117,eth1,VFPT_L10.24.14.136,0/1
DC-GW1, ,EXT1,10.24.14.191,0/1,lag_4
DC-GW1, ,EXT1,10.24.14.192,0/1,lag_4
[efa-bl-pairs]
Bleaf-1,10.2.2.2,10.2.2.3
```

The following fields are available for the `efa-bl-pairs` section:

Field	Description
Name	Border leaf pair name for reference
Leaf1-IP	IP address of border leaf1 which is part of fabric border leaf nodes.
Leaf2-IP	IP address of border leaf2 which is part of fabric border leaf nodes.

#### Deleting border leaf pair

```
openstack network efa-bl-pair-map delete <id>
```

#### Example

```
openstack network efa-bl-pair-map delete 0d19abc5-d31b-4682-b86d-cc1617ea0c97
```

#### Listing border leaf pairs

```
openstack network efa-bl-pair-map list <id>
```

#### Example

```
openstack network efa-bl-pair-map list
+-----+-----+-----+
| ID | Name | Switches |
+-----+-----+-----+
| 0d19abc5-d31b-4682-b86d-cc1617ea0c97 | blp1 | 10.2.2.2,10.2.2.3 |
+-----+-----+-----+
```

#### Neutron EFA console commands

For convenience, Neutron EFA console commands are also provided in the topology plugin for adding, removing, and listing border leaf pairs.

For more details on the commands in this section, see [EFA OpenStack Service Command Reference](#) on page 65..

### Creating border leaf pairs

```
efa-blpair-mapping add <id>
```

#### Example

```
efa-blpair-mapping add --name blp2 --switch-ips 10.2.2.4,10.2.2.3
```

### Removing border leaf pairs

```
efa-blpair-mapping remove <id>
```

#### Example

```
efa-blpair-mapping remove --name blp2
```

### Listing border leaf pairs

```
efa-blpair-mapping list <id>
```

#### Example

```
efa-blpair-mapping list
+-----+-----+-----+
| ID                | Name | Switches          |
+-----+-----+-----+
| 8f3d786b-70e3-4c8e-888c-4f95adc162f4 | blp2 | 10.2.2.4,10.2.2.3 |
+-----+-----+-----+
```

## Journaling

The journal maintains a list of operations that occurred on Neutron and should be mirrored to the controller.

The journal entry is recorded in the pre-commit phase (when applicable) so that when a commit failure occurs, the journal entry and the original operation are deleted.

The journal is processed mainly by a journaling thread that runs periodically and checks whether the journal table has any entries that need to be processed. The thread is triggered in the post-commit phase of the operation (when applicable).

Take **create network** as an example. After it is stored in the Neutron DB by the ML2 plugin, the EFA mechanism driver stores a journal entry representing that operation and triggers the journaling thread to take care of the entry.



#### Note

Journaling is available for the ML2 plugin and Layer 3 plugin.

The Journaling feature delivers additional capabilities:

- EFA reachability and early failure of OpenStack commands.
- A periodic thread updates the reachability of EFA in the Neutron database. This periodic thread runs every 150 seconds, by default, and it is user configurable via the `efa_status_interval` field in

the `ml2_extreme_conf` file. If EFA reachability is down, then the driver does not queue the operations into the journal and the corresponding operation is returned as failed.

- EFA errors are classified as non-retriable and retriable, based on the error code. If the errors are non-retriable, the journal entries are marked as FAILED immediately, without retrying. Currently, except for the connectivity and authorization issue, all other EFA errors are considered as non-retriable.
- If the parent entry moves to FAILED state, then all the dependent entries will also be marked as FAILED there by avoiding unnecessary processing.
- Dependencies for a new journal entry are created at the time of recording with the pending and processing entries in the journal. While a journal entry is being processed, if the parent entry failed, the child entry is made to fail immediately, without trying to process.

For more information, see [https://docs.openstack.org/networking-odl/latest/contributor/drivers\\_architecture.html](https://docs.openstack.org/networking-odl/latest/contributor/drivers_architecture.html).

Several commands are available for journaling.

- **efa-journal list**
- **efa-journal reset**
- **efa-journal clear**

For more information, see [EFA OpenStack Service Command Reference](#) on page 65.

## Sync EFA with Neutron

The OpenStack integration enables you to sync all entries from Neutron to EFA.

### About This Task

When EFA goes out of sync with Neutron configuration, run **efa-sync** to sync all entries from Neutron to EFA. For example, if there is a network and associated ports that are present on Neutron but are not present on EFA, running **efa-sync** creates the entities on EFA. Similarly, if there are stale entries on EFA, they are cleaned up.

You can display summary or detailed information about entities that are out of sync between Neutron and EFA configurations.

### Procedure

1. Sync EFA with Neutron.

```
$ efa-sync execute
Starting Sync
  Syncing Networks..
    Add Network to EFA Id: 1fb97c4a-1a61-4fe4-8ddf-d832f4d37e1e (Succeeded)
  Syncing Subnets..
    Add Subnet to EFA Id: 318ccc16-97ee-4e82-910f-4eee858f25d7 (Succeeded)
  Syncing Ports..
    Add Port to EFA Id: 1fb97c4a-1a61-4fe4-8ddf-d832f4d37e1e-10.20.54.68-port-channel-lag_ext1 (Succeeded)
  Syncing Routers..
    Add Router to EFA Id: e31dec93-10b2-4f75-bdae-da8a0a882475 (Succeeded)
  Syncing Router Ports..
    Add Router Port to EFA SubnetId 318ccc16-97ee-4e82-910f-4eee858f25d7 RouterId e31dec93-10b2-4f75-bdae-da8a0a882475 (Succeeded)
  Summary:
    Neutron Networks added to EFA : 1
```

```

Neutron Networks deleted from EFA : 0
Neutron Ports added to EFA : 1
Neutron Ports deleted from EFA : 0
Neutron Subnets to be added to EFA : 1
Neutron Subnets to be deleted from EFA : 0
Neutron Routers to be added to EFA : 1
Neutron Routers to be deleted from EFA : 0
Neutron Router Ports to be added to EFA : 1
Neutron Router Ports to be deleted from EFA : 0
Neutron Router Routes to be added to EFA : 0
Neutron Router Routes to be deleted from EFA : 0
Total # of resources resynced : 5
Total # of resources failed resync : 0
Clearing failed journal entries...
Number of rows cleared from journal: 0
Completed Sync

```

2. Display summary information about entities that are out of sync between Neutron and EFA configurations.

```

$ efa-sync check-summary
05-November-2020 11:04:15 : Starting Check
Summary:
    Neutron Networks to be added to EFA : 1
    Neutron Networks to be deleted from EFA : 0
    Neutron Ports to be added to EFA : 2
    Neutron Ports to be deleted from EFA : 0
    Total # of resources to be resynced : 3
05-November-2020 11:04:16 : Completed Check

```

Journal entries that are in a failed state are cleared from the journal list if `efa-sync` is successful. If the `efa-sync` run fails (complete failure or partial failure), the failed journal entries are not cleared.

3. Display summary or detailed information about those entities that are out of sync between Neutron and EFA configurations.

```

$ efa-sync check-summary
05-November-2020 11:04:15 : Starting Check
Summary:
    Neutron Networks to be added to EFA : 1
    Neutron Networks to be deleted from EFA : 0
    Neutron Ports to be added to EFA : 2
    Neutron Ports to be deleted from EFA : 0
    Total # of resources to be resynced : 3
05-November-2020 11:04:16 : Completed Check

$ efa-sync check-detail
05-November-2020 11:04:22 : Starting Check
    Neutron Networks to be added to EFA :
        Id: 9e63bc57-568f-488e-9214-21db3fd3cd12
    Neutron Ports to be added from EFA :
        Id: 9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.11-port-channel-
lag_1
        Id: 9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.46-port-channel-
lag_1
05-November-2020 11:04:22 : Completed Check

```



#### Note

This tool logs the output of the console to the `efa-sync-console.log` file.



## Verify EFA Health

The `efa-health show` command is provided as part of the installation package.

You can use the command to verify the following: connectivity with EFA, database write operations, OpenStack journal thread status, and the status of the EFA Kubernetes cluster.



### Tip

This feature creates a VRF for testing the health status. So the tenant `vrf-count` option accounts for this additional VRF.

In the command output, the **EFA Health** entry represents EFA reachability and EFA database health.

In the output, the **EFA High Availability Status** entry shows DOWN even if only one node is DOWN. However, even with one node down, EFA Health remains UP because EFA is still operational.

Run `efa-health show` to show a summary of health status.

In this output, EFA health is the consolidated health of both EFA reachability and EFA DB health. EFA high availability status will be DOWN even if one of the nodes is DOWN but EFA health can still be UP as the EFA is still operational.

```
$ sudo efa-health show
EFA Host                : efa.extremenetworks.com
EFA Health              : UP
EFA Health Failure Reason :
EFA High Availability Status : UP
EFA Master              : UP
EFA Slave               : UP

EFA Neutron Journal Size within Threshold : NO
EFA Neutron Journal Size Threshold Setting : 5
EFA Neutron Current Journal Size         : 7
EFA Neutron Journal Maintenance Status   : UP
```

Run `efa-health show` with the `--advanced` option, which can give EFA POD or cluster status, in detail. This can be used to infer the pods running status on EFA.

```
$ sudo efa-health show --advanced
EFA Host                : efa.extremenetworks.com
EFA Health              : UP
EFA Health Failure Reason :
EFA High Availability Status : UP
EFA Master              : UP
EFA Slave               : UP

EFA Neutron Journal Size within Threshold : NO
EFA Neutron Journal Size Threshold Setting : 5
EFA Neutron Current Journal Size         : 7
EFA Neutron Journal Maintenance Status   : UP

EFA Cluster Status      : 18/28 Pods Running
  rabbitmq-0            : Running
  efa-api-docs-67b8c76ddb-2hg7s      : Running
  gosnmp-service-546d76b6f9-7wb6t    : Running
  goopenstack-service-5b687f9f8c-srlxx : Running
  gosystem-service-76bbbc6d-jzkng    : Running
  goauth-service-75c88f4986-vqppm    : Running
  goraslog-service-7467fb7759-jm28m  : Running
```

```

gotenant-service-5f8bc9f458-zbz21      : Running
gohyperv-service-b546d647f-ws7t8      : Running
gorbac-service-b546dbdd5-v9dml        : Running
gofabric-service-65b6f7d4cf-rfq6f     : Running
gonotification-service-5fb74fc959-2x9lt : Running
goinventory-service-79d6545c69-s66tx   : Running
govcenter-service-7c6cb944dd-2d81t    : Running
goinventory-service-79d6545c69-gxtn6   : Init:0/2
efa-api-docs-67b8c76ddb-n8zfp        : Running
gosnmp-service-546d76b6f9-fv8tz       : Init:0/1
goraslog-service-7467fb7759-rvppq     : Running
gorbac-service-b546dbdd5-s9qhx        : Init:1/2
gosystem-service-76bbbc6d-x8x8g      : Running
govcenter-service-7c6cb944dd-rx7gd    : Init:0/2
gohyperv-service-b546d647f-zjq97     : Init:0/2
goauth-service-75c88f4986-9fbn7      : Init:0/2
gofabric-service-65b6f7d4cf-q7qns    : Init:0/2
gonotification-service-5fb74fc959-rr724 : Init:0/2
gotenant-service-5f8bc9f458-n7gm8     : Init:0/3
goopenstack-service-5b687f9f8c-bfssf  : Init:0/2
rabbitmq-1                             : Running

```

## Syslog Integration

Syslog configuration enables you to consolidate multiple controller logs in the same place.

OpenStack generates the following log files:

- `neutron-server.log`
- `efa-blpair-mapping.log`
- `efa-health.log`
- `efa-journal.log`
- `efa-link-mapping.log`
- `efa-sync.log`

Logs in OpenStack are output to the `/var/log/neutron` directory within the controller, outside the docker shell.

In case of multiple controllers, these files are separately available in each controller. Syslog integration is possible by setting the appropriate rules and syslog server settings. See syslog documentation for more information.



### Note

The configuration files for rsyslog must be updated with the appropriate logging rules for the logs to be available at `/var/log/neutron`. By default, the application log files are available in the `/var/log/catch-all/` directory.


## Multiple VIM/VPOD Instances

Each VIM/VPOD instance or Open Stack is mapped to a separate EFA tenant. A VIM instance can be segregated as follows:

- Grouping all physical interfaces together through topology specification

- Restricting the VLAN range details in the Neutron initialization file

The following figure shows an example of multiple VIM/VPOD instances managing the same IP fabric.



**Figure 13: Multiple VIM or VPOD instances**

VIM	EFA Tenant	BD	Properties
VIM1	OS-Tenant-VIM1	BD is disabled. CTAG is mapped to VLAN on the IP fabric.	VLAN-Range = 2-4080 specified in ml2_conf.ini network_vlan_ranges = physnet1:2:4080 Physical interfaces can be specified using the topology.
VIM2	OS-Tenant-VIM2	BD is enabled. CTAG is mapped to BD on the IP fabric.	LAN-Range = 2-4080 specified in ml2_conf.ini network_vlan_ranges = physnet1:2:4080 Tenant created with <b>-bd-enable flag</b> . Physical interfaces can be specified using the topology.

The following example shows multiple VIM/VPOD instances.

```
# efa tenant create --name VIM1 --port 10.24.80.111[0/1],10.24.80.112[0/1] --vlan-range 2-4080 --vrf-count 200

# efa tenant create --name VIM2 --port 10.24.80.112[0/2],10.24.80.113[0/1],10.24.80.114[0/1] --enable-bd --vlan-range 2-4080 --vrf-count 200

# efa tenant show
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Name | L2VNI-Start | L3VNI-Start | VLAN-Range | VRF-Count | Enable-BD |
+-----+-----+-----+-----+-----+-----+
```

```

|          Ports          |
+-----+-----+-----+-----+-----+
| default-tenant | *          | *          | *          | *          | *          |
*                  |           |           |           |           |           |
+-----+-----+-----+-----+-----+
| VIM1           | 0          | 0          | 2-4080     | 200        | False      |
10.24.80.111[0/1] |           |           |           |           |           |
|               |           |           |           |           |           |
10.24.80.112[0/1] |           |           |           |           |           |
+-----+-----+-----+-----+-----+
| VIM2           | 0          | 0          | 2-4080     | 200        | True       |
10.24.80.112[0/2] |           |           |           |           |           |
|               |           |           |           |           |           |
10.24.80.113[0/1] |           |           |           |           |           |
|               |           |           |           |           |           |
10.24.80.114[0/1] |           |           |           |           |           |
+-----+-----+-----+-----+-----+
*: VNIs will be allocated from the default-tenant VNI pool, for new tenant and default-
tenant network.
--- Time Elapsed: 23.9927ms ---

```

## Scale-in and Scale-out of Compute Nodes

**Table 15: Scale-in of compute nodes**

Scale-in	EFA Impact
Link removal from default physical network	VLAN provisioning for endpoints is done when Neutron ports are unbound to a host or compute. Hence, removing a link which belongs to the default physical network does not have any impact, only the link-mapping table is updated.
Link removal from non-default physical networks	For non-default physical networks, VLAN provisioning is done at the network level. Hence, the links that are removed from non-default networks are automatically removed from the existing networks. On EFA, the interface is removed from the corresponding EPGs. If there are any errors during EFA unbinding, the operation succeeds. This behavior is consistent with the delete port operation.
Links bound to port-channels on non-default physical networks	Removing a link which maps to a LAG interface does not affect EPG unless the last member of the LAG is removed in link mapping. For example, LAGs have two entries in the link mapping. Removing only one entry does not affect the existing EPG. Only when the last entry for that LAG is removed, EFA removes the LAGs from corresponding EPGs.

**Table 16: Scale-out of compute nodes**

Scale-out	EFA Impact
Link addition to default physical network	VLAN provisioning of endpoints is done when Neutron ports are bound to a host or compute. Hence, adding a link which belongs to the default physical network does not have any impact, only the link-mapping table is updated.
Link addition to non-default physical networks	For non-default physical networks, VLAN provisioning is done during network-create time. Hence, the new links that are added to non-default networks are automatically added to the existing networks. On EFA, the interface is added to the corresponding EPGs. If there are any errors during EFA binding, the whole operation fails.

**Table 16: Scale-out of compute nodes (continued)**

Scale-out	EFA Impact
Links bound to port-channels on non-default physical networks	If the switch link is a LAG interface, as soon as the first link to LAG is added, the corresponding LAG is added to the EPGs. Further link addition on that LAG requires no further changes for EFA.
File option in efa-link-add	File option enables saving link mappings in a file for bulk configuration. If there is an error in adding a link during file replay, the operation is aborted. You can correct the errors and replay it.

## Virtual Machine Migration

Virtual Machine Migration (VMotion) enables migration of live virtual machines from one OpenStack compute server to another. You can use VMotion to migrate VMs during planned maintenance or redistribute load on the server.

In non-live or cold migration, the VM instances are shut down before migrating them to another server, resulting in disruption of services. In live migration, the VM instances continue to run during migration without disrupting any services.

### Enable VMotion

#### Procedure

1. Set the following parameters in `nova.conf` on all compute nodes. Ensure that `instances_path` and `restorecon` are same for all compute nodes.



#### Note

This setting allows VNC clients from any IP address to connect to instance consoles. Ensure to take additional measures to secure networks.

```
vncserver_listen = 0.0.0.0
instances_path = /var/lib/nova/instances
restorecon = /etc/hosts
```

2. Enable password-less SSH.
 

The libvirt daemon that runs as root uses SSH protocol to copy the instance to the destination.
3. Configure the firewalls to allow libvirt to communicate with compute nodes.
 

The default libvirt TCP port range is 49152 to 49261 for copying memory and disk contents.

### Migrate Virtual Machines

#### Procedure

1. View the OpenStack server list to determine the VM to be migrated.

```
# openstack server list
```

```

+-----+-----+-----+-----+-----+
| ID | Name | Status | Networks | Image Name |
+-----+-----+-----+-----+-----+
| d1df1b5a-70c4-4fed-98b7-423362f2c47c | vm1 | ACTIVE | private=a.b.c.d | ... |
| d693db9e-a7cf-45ef-a7c9-b3ecb5f22645 | vm2 | ACTIVE | private=e.f.g.h | ... |
+-----+-----+-----+-----+-----+

```

2. Select the destination host.
  - For manual selection of the destination host, proceed to the next step.
  - For automatic selection of the destination host, go to Step 6.
3. View the server details of the selected VM.

```

# openstack server show d1df1b5a-70c4-4fed-98b7-423362f2c47c
+-----+-----+
| Field | Value |
+-----+-----+
| ... | ... |
| OS-EXT-SRV-ATTR:host | HostB |
| ... | ... |
| addresses | a.b.c.d |
| flavor | ml.tiny |
| id | d1df1b5a-70c4-4fed-98b7-423362f2c47c |
| name | vm1 |
| status | ACTIVE |
| ... | ... |
+-----+-----+

```

4. Determine the destination host to migrate the selected VM.

```

# openstack compute service list
+-----+-----+-----+-----+-----+-----+
+-----+
| ID | Binary | Host | Zone | Status | State | Updated |
+-----+-----+-----+-----+-----+-----+
| 3 | nova-conductor | HostA | internal | enabled | up | 2017-02-18T09:42:29.000000 |
| 4 | nova-scheduler | HostA | internal | enabled | up | 2017-02-18T09:42:26.000000 |
| 5 | nova-consoleauth | HostA | internal | enabled | up | 2017-02-18T09:42:29.000000 |
| 6 | nova-compute | HostB | nova | enabled | up | 2017-02-18T09:42:29.000000 |
| 7 | nova-compute | HostC | nova | enabled | up | 2017-02-18T09:42:29.000000 |
+-----+-----+-----+-----+-----+-----+
+-----+

# openstack host show HostC
+-----+-----+-----+-----+-----+
| Host | Project | CPU | Memory MB | Disk GB |
+-----+-----+-----+-----+-----+
| HostC | (total) | 16 | 32232 | 878 |
| HostC | (used_now) | 22 | 21284 | 422 |
| HostC | (used_max) | 22 | 21284 | 422 |
| HostC | p1 | 22 | 21284 | 422 |
| HostC | p2 | 22 | 21284 | 422 |
+-----+-----+-----+-----+-----+

```

5. Migrate the VM instance.

- Manual selection of the destination host:

```
# openstack server migrate d1df1b5a-70c4-4fed-98b7-423362f2c47c --live HostC
```

- Automatic selection of the destination host:

```
# nova live-migration d1df1b5a-70c4-4fed-98b7-423362f2c47c
```

6. View the host server details of the migrated VM to confirm the status of migration. If migration fails, go to Step 8.

```
# openstack server show d1df1b5a-70c4-4fed-98b7-423362f2c47c
```

Field	Value
...	...
OS-EXT-SRV-ATTR:host	<b>HostC</b>
...	...

7. (Optional) View the log files on the controller and compute nodes for more information about migration failure.

- nova-scheduler
- nova-conductor
- nova-compute

8. (Optional) Stop the migration manually.

```
# nova live-migration-abort INSTANCE_ID MIGRATION_ID
```

9. (Optional) Force complete the migration.

```
# nova live-migration-force-complete INSTANCE_ID MIGRATION_ID
```

## Add Certificate to OpenStack Controller

### Procedure

1. Copy certificate from EFA to OpenStack Controller.

```
administrator@osController-55:$ su
Password:
root@osController-55:# cd /usr/local/share/ca-certificates
root@osController-55:#
root@osController-55:/usr/local/share/ca-certificates# scp root@10.21.88.130://usr/
local/share/ca-certificates/extreme-ca-chain.crt .
root@10.21.88.130's password:
extreme-ca-chain.crt
100% 4488    4.4KB/s   00:00
root@osController-55:/usr/local/share/ca-certificates#
```

2. Restart Neutron to apply the token and certificate.

```
# sudo service neutron-* restart
```





# EFA OpenStack Service Command Reference

---

- [efa openstack debug](#) on page 66
- [efa openstack debug router-route delete](#) on page 67
- [efa openstack execution](#) on page 68
- [efa openstack network show](#) on page 69
- [efa openstack network-interface show](#) on page 70
- [efa openstack router show](#) on page 71
- [efa openstack router-interface show](#) on page 72
- [efa openstack router-route show](#) on page 72
- [efa openstack subnet show](#) on page 73
- [efa openstack sync start](#) on page 74
- [efa-blpair-mapping](#) on page 74
- [efa-health show](#) on page 76
- [efa-journal clear](#) on page 78
- [efa-journal list](#) on page 79
- [efa-journal list deps](#) on page 79
- [efa-journal reset](#) on page 81
- [efa-sync execute](#) on page 82
- [openstack network efa-bl-pair-map create](#) on page 84
- [openstack network efa-topology-link-map create](#) on page 85
- [openstack network efa-topology-link-map delete](#) on page 86
- [openstack network efa-topology-link-map list](#) on page 87
- [openstack router create](#) on page 88

This chapter describes OpenStack commands issued from EFA.

## efa openstack debug

---

Displays OpenStack debug information.

### Syntax

```
efa openstack debug [ network | network-interface | tenant | router | router-interface ]
```

```
efa openstack debug network delete --neutron-id
```

Deletes the network and its summary information.

```
efa openstack debug network-interface delete --neutron-id
```

Deletes the network interface and its summary information.

```
efa openstack debug router delete --router-id
```

Deletes the router and its summary information.

```
efa openstack debug router-interface delete --router-id <router-id> --subnet-id <subnet-id>
```

Deletes the router interface and its summary information.

```
efa openstack debug tenant cleanup --name tenant name
```

Deletes a network.

### Parameters

**cleanup** --*name*

Cleans up all OpenStack assets associated to a tenant

**delete** --*neutron-id* | --*router-id* | *subnet-id*

Deletes the selected network element

**network**

Specifies network commands

**network-interface**

Specifies network interface commands

**tenant**

Specifies tenant commands

**router**

Specifies router commands

**router-interface**

Specifies router interface commands

**neutron-id** *string*

Neutron ID of the network

**router-id** *string*  
Router ID

**subnet-id** *string*  
Subnet ID

**name** *string*  
Tenant name

**help** *string*  
Displays help on the command

## efa openstack debug router-route delete

---

Debug command for deleting routes in the specified router.

```
efa openstack debug router-route delete [--router-id router-id | --  
destination destination | --nexthop ip]
```

### Parameters

**--router-id** *router-id*  
Specifies a comma-separated range of router IDs.

**--destination** *destination*  
Submask

**--nexthop** *ip*  
IP

For IPv4:

```
efa openstack debug router-route delete --router-id f8a44de0-fc8e-45df-93c7-f79bf3b01c95  
--destination 10.56.6.0/22 --nexthop 10.56.6.1
```

For IPv6:

```
efa openstack debug router-route delete  
--router-id f8a44de0-fc8e-45df-93c7-f79bf3b01c95 --destination 101::/64 --nexthop 11::1
```

## efa openstack execution

---

Provides OpenStack execution commands.

### Syntax

```
efa openstack execution delete [ --days int32 | --help ]
```

```
efa openstack execution show [ --help | --id | --limit int32 | --status ]
```

### Parameters

#### **delete**

Deletes execution entries older than the specified days

#### **--days int32**

Deletes execution entries older than the specified days. Default is 30.

#### **--id**

Filters the executions based on execution id. "limit" and "status" flags are ignored when "id" flag is given.

#### **--help**

Provides help for execution

#### **--limit int32**

Limits the number of executions to be listed. Value "0" will list all the executions. Default is 10.

#### **show**

Displays the list of executions.

#### **--status**

Filters the executions based on the status (failed, succeeded, all). Default is all.

## efa openstack network show

---

Displays OpenStack network information.

### Syntax

```
efa openstack network show
```

### Examples

This example shows typical results.

```
efa openstack network show
+-----+-----+-----+-----+
|           Neutron ID           | Tenant | MTU | CTAG |
+-----+-----+-----+-----+
| 9237b75d-6fa9-404b-8fab-979b7d384c40 | RegionOne | 2100 | 3 |
+-----+-----+-----+-----+
| 23313c80-0b91-455a-8dce-91e6cb7424a7 | RegionOne | 2100 | 2 |
+-----+-----+-----+-----+
```

## efa openstack network-interface show

Displays OpenStack network-interface information.

### Syntax

```
efa openstack network-interface show
```

### Examples

This example shows typical results.

```
# efa openstack network-interface show

+-----+-----+-----+
+-----+-----+-----+
|           Neutron Port ID           |           Neutron Network ID           | Switch
IP   | Switch Interface |           |           |
+-----+-----+-----+
+-----+-----+-----+
| 123e4567-e89b-12d3-a456-426655440001 | 123e4567-e89b-12d3-a456-426655440001 |
10.24.80.134 | 0/9           |
+-----+-----+-----+
+-----+-----+-----+
| 123e4567-e89b-12d3-a456-426655440003 | 123e4567-e89b-12d3-a456-426655440001 |
10.24.80.133 | 0/9           |
+-----+-----+-----+
+-----+-----+-----+
| 123e4567-e89b-12d3-a456-426655440005 | 123e4567-e89b-12d3-a456-426655440002 |
10.24.80.134 | 0/9           |
+-----+-----+-----+
+-----+-----+-----+
| 123e4567-e89b-12d3-a456-426655440007 | 123e4567-e89b-12d3-a456-426655440002 |
10.24.80.133 | 0/9           |
+-----+-----+-----+
+-----+-----+-----+
```

## efa openstack router show

Displays OpenStack router information.

### Syntax

```
efa openstack router show
```

### Examples

This example shows typical results.

```
efa openstack router show
+-----+-----+-----+
+-----+
|           Router ID           |           VRF Name           | Tenant |
Routing type |
+-----+-----+-----+
+-----+
| bc482e46-8896-442c-a29d-2dd76abca2ae | bc482e468896442ca29d2dd76abca2ae | RegionOne |
centralized |
+-----+-----+-----+
+-----+
| 99ec4fd8-fbc4-4117-a529-5e9a40049906 | 99ec4fd8fbc44117a5295e9a40049906 | RegionOne |
distributed |
+-----+-----+-----+
+-----+
```

## efa openstack router-interface show

Displays OpenStack router-interface information.

### Syntax

```
efa openstack router-interface show
```

### Examples

This example shows typical results.

```
# efa openstack router-interface show

+-----+-----+
|          Subnet ID          |          Router ID          |
+-----+-----+
| 323e4567-e89b-12d3-a456-426655440001 | 523e4567-e89b-12d3-a456-426655440001 |
+-----+-----+
| 323e4567-e89b-12d3-a456-426655440002 | 523e4567-e89b-12d3-a456-426655440001 |
+-----+-----+
```

## efa openstack router-route show

Shows routes on the specified router.

```
efa openstack router-route show [--id router-id]
```

### Parameters

**--id** *router-id*

Router ID

```
efa openstack router-route show --id f8a44de0-fc8e-45df-93c7-f79bf3b01c95

+-----+-----+-----+-----+
|          Router ID          | Destination | Nexthop IP | Device |
+-----+-----+-----+-----+
| f8a44de0-fc8e-45df-93c7-f79bf3b01c95 | 10.56.6.0/22 | 10.56.6.1 |      |
| 10.25.25.100,10.25.25.100 |             |             |      |
+-----+-----+-----+-----+
| f8a44de0-fc8e-45df-93c7-f79bf3b01c95 | 10.56.6.0/22 | 10.56.6.1 |      |
| 10.25.25.100,10.25.25.101 |             |             |      |
+-----+-----+-----+-----+
| f8a44de0-fc8e-45df-93c7-f79bf3b01c95 | 101::/64    | 11:::1    |      |
| 10.25.25.100,10.25.25.101 |             |             |      |
+-----+-----+-----+-----+
| f8a44de0-fc8e-45df-93c7-f79bf3b01c95 | 101::/64    | 11:::1    |      |
| 10.25.25.100,10.25.25.101 |             |             |      |
+-----+-----+-----+-----+
```



## efa openstack subnet show

Displays OpenStack subnet information.

### Syntax

```
efa openstack subnet show
```

### Examples

This example shows typical results.

```
efa openstack subnet show
+-----+-----+-----+-----+
+-----+-----+-----+-----+
|          Subnet ID          |          Network ID          |
| CIDR      | Gateway IP | IPv6RAMode | IPv6AddressMode |
+-----+-----+-----+-----+
| 869dbbcd-2c80-440a-a42f-1e85cab60b98 | 9237b75d-6fa9-404b-8fab-979b7d384c40 |
| 110:2000::/64 | 110:2000::1 | dhcpv6-stateful | |
+-----+-----+-----+-----+
| 063eec5a-dbb3-480b-b1d0-1d7fc4715b67 | 23313c80-0b91-455a-8dce-91e6cb7424a7 |
| 10:2000::/64 | 10:2000::1 | | |
+-----+-----+-----+-----+
| ea233dda-b692-4a46-adc8-ea58ac07f4d2 | 23313c80-0b91-455a-8dce-91e6cb7424a7 |
| 8.8.8.0/24 | 8.8.8.1 | | |
+-----+-----+-----+-----+
| db91323d-fcdd-4450-87b2-57bb7e033c0c | 9237b75d-6fa9-404b-8fab-979b7d384c40 |
| 18.8.8.0/24 | 18.8.8.1 | | |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
```

## efa openstack sync start

---

Syncs entries such as networks, network interfaces, routers, and router interfaces from OpenStack to a tenant if they are out of sync.

### Syntax

```
efa openstack sync start [ --tenant tenant name | --help ]
```



#### Note

There is no output for this command.

### Parameters

#### **--tenant**

The name of the tenant

#### **--help**

Provides help for execution

## efa-blpair-mapping

---

Adds, removes, or lists border leaf pairs.

```
efa-blpair-mapping add [ --name name | --switch-ips switches ]
```

```
efa-blpair-mapping remove [ --name name ]
```

```
efa-blpair-mapping list
```

### Parameters

#### **add**

Adds the specified border leaf pair.

#### **remove**

Removes the specified border leaf pair.

#### **list**

Lists border leaf pairs.

#### **--name** *name*

Specifies the IP address of the device for which you want to see a list of interfaces.

#### **--switch-ips**

Comma separated IP addresses of border-leaf on fabric. For example, 10.2.2.2,10.2.2.3.

The following example adds the specified border leaf pair.

```
$ efa-blpair-mapping add --name blp2 --switch-ips 10.2.2.4,10.2.2.3
```

The following example removes the specified border leaf pair.

```
$ efa-blpair-mapping remove --name blp2
```

The following example lists the configured border leaf pairs.

```
$ efa-blpair-mapping list
+-----+-----+-----+
|          ID          | Name |   Switches   |
+-----+-----+-----+
| 8f3d786b-70e3-4c8e-888c-4f95adc162f4 | blp2 | 10.2.2.4,10.2.2.3 |
+-----+-----+-----+
```

## efa-health show

Verifies the following: connectivity with EFA, database write operations, OpenStack journal thread status, and the status of the EFA Kubernetes cluster.

### Syntax

```
efa-health show [--advanced ]
```

### Parameters

#### **--advanced**

Provides detailed status about EFA pods and the cluster. Omit this parameter to return basic status information.

### Usage Guidelines

In the output, the EFA Health entry represents EFA reachability and EFA database health.

In the output, the EFA High Availability Status entry shows DOWN even if only one node is DOWN. However, even with one node down, EFA Health remains UP because EFA is still operational.

### Examples

This example shows basic output for the command.

```
$ sudo efa-health show
EFA Host                : efa.extremenetworks.com
EFA Health               : UP
EFA Health Failure Reason :
EFA High Availability Status : UP
EFA Master               : UP
EFA Slave                : UP

EFA Neutron Journal Size within Threshold : NO
EFA Neutron Journal Size Threshold Setting : 5
EFA Neutron Current Journal Size         : 7
EFA Neutron Journal Maintenance Status   : UP
```

Run `efa-health show` with the `--advanced` option, which can give EFA POD or cluster status, in detail. This can be used to infer the pods running status on EFA.

```
$ sudo efa-health show --advanced
EFA Host                : efa.extremenetworks.com
EFA Health               : UP
EFA Health Failure Reason :
EFA High Availability Status : UP
EFA Master               : UP
EFA Slave                : UP

EFA Neutron Journal Size within Threshold : NO
EFA Neutron Journal Size Threshold Setting : 5
EFA Neutron Current Journal Size         : 7
EFA Neutron Journal Maintenance Status   : UP

EFA Cluster Status      : 18/28 Pods Running
```

```
rabbitmq-0 : Running
efa-api-docs-67b8c76ddb-2hg7s : Running
gosnmp-service-546d76b6f9-7wb6t : Running
goopenstack-service-5b687f9f8c-srlxx : Running
gosystem-service-76bbbc6d-jzkng : Running
goauth-service-75c88f4986-vqppm : Running
goraslog-service-7467fb7759-jm28m : Running
gotenant-service-5f8bc9f458-zbz21 : Running
gohyperv-service-b546d647f-ws7t8 : Running
gorbac-service-b546dbdd5-v9dml : Running
gofabric-service-65b6f7d4cf-rfq6f : Running
gotification-service-5fb74fc959-2x91t : Running
goinventory-service-79d6545c69-s66tx : Running
govcenter-service-7c6cb944dd-2d81t : Running
goinventory-service-79d6545c69-gxtn6 : Running
efa-api-docs-67b8c76ddb-n8zfp : Running
gosnmp-service-546d76b6f9-fv8tz : Init:0/1
goraslog-service-7467fb7759-rvppq : Running
gorbac-service-b546dbdd5-s9qxh : Init:1/2
gosystem-service-76bbbc6d-x8x8g : Running
govcenter-service-7c6cb944dd-rx7gd : Init:0/2
gohyperv-service-b546d647f-zjq97 : Init:0/2
goauth-service-75c88f4986-9fbn7 : Init:0/2
gofabric-service-65b6f7d4cf-q7qns : Init:0/2
gotification-service-5fb74fc959-rr724 : Init:0/2
gotenant-service-5f8bc9f458-n7gm8 : Init:0/3
goopenstack-service-5b687f9f8c-bfssf : Init:0/2
rabbitmq-1 : Running
```

## efa-journal clear

---

Removes all entries or entries specified according to state from the journal.

### Syntax

```
efa-journal clear [ pending | failed | completed | processing | all ]
```

### Parameters

**pending**

Clears all entries in Pending state.

**failed**

Clears all entries in Failed state.

**completed**

Clears all entries in Completed state.

**processing**

Clears all entries in Processing state.

**all**

Clears all entries.

### Usage Guidelines

**Important**

Stop the Neutron service before clearing entries using the **all**, **pending**, **failed**, or **processing** parameters. Otherwise, the processing thread can repopulate the entries in the journal.

### Examples

The following example clears all entries in the Failed state.

```
$ efa-journal clear failed
```

The following example clears all entries. Note the warning to stop the Neutron service.

```
$ efa-journal clear all
Warning!!! Clearing states pending/processing/all should be executed after shutting
down neutron service. Otherwise the entries can be re-populated by the processing thread.
```

The following example clears all entries in Processing state.

```
$ efa-journal clear processing
Warning!!! Clearing states pending/processing/all should be executed after shutting
down neutron service.
Otherwise the entries can be re-populated by the processing thread.
```

The following example clears all entries in Pending state.

```
$ efa-journal clear pending
Warning!!! Clearing states pending/processing/all should be executed after shutting down
neutron service.
Otherwise the entries can be re-populated by the processing thread.
```

## efa-journal list

Lists all journal entries that are present in the extreme journaling database, along with their statuses.

### efa-journal list

This example lists all journal entries and their statuses.

```
$ efa-journal list
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Id | Object Type | Object UUID | Operation | state |
Retry Count | Reason |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 77 | network | 208421ad-483e-4e45-8c6f-d135146a1ee7 | create | failed |
5 | EFA connection timed out |
| 78 | subnet | 0443d876-eafb-4b8c-9c83-7d32c2cccad6 | create | failed |
0 | Parent Journal ID: 77 failed |
| 79 | port | 9cb6d418-d916-4be8-aad5-1cb07d388d5f | update | failed |
0 | Parent Journal ID: 77 failed |
| 80 | network | b22415ef-afdc-46fa-834f-8506c132da08 | create | failed |
5 | EFA connection timed out |
| 81 | subnet | 2a1e3aea-c996-45f1-94f9-84b604597dc3 | create | failed |
0 | Parent Journal ID: 80 failed |
| 82 | port | 4f3fe359-8de9-45f6-be4b-b9d7c565399f | update | failed |
0 | Parent Journal ID: 80 failed |
| 83 | network | 2237342e-95ec-44c2-92e6-706092bleda3 | create | processing |
0 | |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
```

## efa-journal list deps

Debug command: dumps the journal dependency table in user-readable format.

### efa-journal list deps

This example dumps the journal dependency table.

```
$ efa-journal list deps
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Parent | Parent-obj | Parent-uuid | Parent-op | Parent-state |
Child | Child-obj | Child-uuid | Child-op | Child-state |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 77 | network | 208421ad-483e-4e45-8c6f-d135146a1ee7 | create | failed |
| 78 | subnet | 0443d876-eafb-4b8c-9c83-7d32c2cccad6 | create | failed |
| 77 | network | 208421ad-483e-4e45-8c6f-d135146a1ee7 | create | failed |
| 79 | port | 9cb6d418-d916-4be8-aad5-1cb07d388d5f | update | failed |
| 78 | subnet | 0443d876-eafb-4b8c-9c83-7d32c2cccad6 | create | failed |
| 79 | port | 9cb6d418-d916-4be8-aad5-1cb07d388d5f | update | failed |
| 80 | network | b22415ef-afdc-46fa-834f-8506c132da08 | create | failed |
| 81 | subnet | 2a1e3aea-c996-45f1-94f9-84b604597dc3 | create | failed |
```

	80		network		b22415ef-afdc-46fa-834f-8506c132da08		create		failed
	82		port		4f3fe359-8de9-45f6-be4b-b9d7c565399f		update		failed
	81		subnet		2a1e3aea-c996-45f1-94f9-84b604597dc3		create		failed
	82		port		4f3fe359-8de9-45f6-be4b-b9d7c565399f		update		failed
	83		network		2237342e-95ec-44c2-92e6-706092b1eda3		create		failed
	84		subnet		51c32ef1-c15e-4a1a-ad89-6e2adfdf8da5		create		failed
	83		network		2237342e-95ec-44c2-92e6-706092b1eda3		create		failed
	85		port		aab84828-5e08-4a95-aef6-2498f6ef5991		update		failed
	84		subnet		51c32ef1-c15e-4a1a-ad89-6e2adfdf8da5		create		failed
	85		port		aab84828-5e08-4a95-aef6-2498f6ef5991		update		failed
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+									



## efa-journal reset

---

Resets the retry count for all or specified failed journal entries and retries the operation.

### Syntax

```
efa-journal reset [ Journal_ID ]
```

### Parameters

*Journal\_ID*

Specifies the journal ID that you want to reset.

### Usage Guidelines

You can return journal IDs by running the `efa-journal list` command.

### Examples

The following example resets specific failed entries - 77 and 78.

```
$ efa-journal reset 77 78  
5dc0bf8b-674c-4c9f-944e-dc048bd184f0
```

The following example resets all failed entries.

```
$ efa-journal reset
```

## efa-sync execute

Syncs EFA with Neutron. Displays summary or detailed information about entities that are out of sync between Neutron and EFA configurations.

### Syntax

**efa-sync execute**

**efa-sync check-summary**

**efa-sync check-detail**

### Usage Guidelines

The command also logs console output into efa-sync-console.log.

The efa-sync tool uses Neutron APIs and keystone authentication. The authentication parameters are selected from the neutron.conf file. Before running efa-sync, make sure that all the parameters under the [keystone-authtoken] section are set to the correct values.



#### Note

In case of default OpenStack settings in neutron.conf, the following parameters need to be edited to 'default' -- lowercase 'd' under [keystone\_authtoken]:

```
project_domain_name = default
user_domain_name = default
```

Journal entries that are completed or failed are cleared from the journal list if the execution of efa-sync is successful. In case of failure, the journal entries are not cleared.

### Examples

This example syncs EFA with Neutron.

```
$ efa-sync execute
Starting Sync
Syncing Networks..
  Add Network to EFA Id: 9e63bc57-568f-488e-9214-21db3fd3cd12
  (Succeeded)
Syncing Ports..
  Add Port to EFA Id:
  9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.11-port-channel-lag_1
  (Succeeded)
  Add Port to EFA Id:
  9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.46-port-channel-lag_1
  (Succeeded)
Completed Sync
```

This example displays summary information about entities that are out of sync between Neutron and EFA configurations.

```
$ efa-sync check-summary
05-November-2020 11:04:15 : Starting Check
Summary:
  Neutron Networks to be added to EFA : 1
```

```
Neutron Networks to be deleted from EFA : 0
Neutron Ports to be added to EFA : 2
Neutron Ports to be deleted from EFA : 0
Total # of resources to be resynced : 3
05-November-2020 11:04:16 : Completed Check
```

This example displays detailed information about entities that are out of sync between Neutron and EFA configurations.

```
$ efa-sync check-detail
05-November-2020 11:04:22 : Starting Check
  Neutron Networks to be added to EFA :
    Id: 9e63bc57-568f-488e-9214-21db3fd3cd12
  Neutron Ports to be added from EFA :
    Id: 9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.11-port-channel-
lag_1
    Id: 9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.46-port-channel-lag_1
05-November-2020 11:04:22 : Completed Check
```

## openstack network efa-bl-pair-map create

Maps a border-leaf pair and specifies the border-leaf pair on which the route will be installed.

### Syntax

```
openstack network efa-bl-pair-map create [--name name | --switch-ips  
    switches]
```

### Parameters

**--name** *name*

Specifies the name of the border-leaf pair

**--switch-ips** *switches*

Comma separated IP addresses of border-leaf on fabric. For example, 10.2.2.2,10.2.2.3.

### Examples

```
openstack network efa-bl-pair-map create --name blp1 --switch-ips 10.2.2.2,10.2.2.3  
+-----+  
| Field      | Value                                     |  
+-----+  
| id         | 0d19abc5-d31b-4682-b86d-cc1617ea0c97 |  
| name       | blp1                                     |  
| switch_ips | 10.2.2.2,10.2.2.3                       |  
+-----+
```

## openstack network efa-topology-link-map create

---

Creates a topology link

### Syntax

```
openstack network efa-topology-link-map create [--host HOST | --nic NIC |  
--provider-network PROVIDER_NETWORK | --port PORT | --po-name PO_NAME  
| --help ]
```

### Parameters

**host**

Host name of the compute node which connects to the switch

**nic**

Physical NIC on the compute host which connects to the switch

**provider-network**

Provider network name on compute host. For example, physnet1 (default).

**port**

Switch port to which the physical NIC on the compute host is connected

**po-name**

Switch PO name where physical ports are aggregated to NIC where compute host is connected.

**help**

Displays help for the command

### Examples

```
# openstack network efa-topology-link-map create --host  
Openstack115 --nic eth1 --pn default --switch 10.24.14.133 --port 0/1 --po-name  
lag_1
```

## openstack network efa-topology-link-map delete

---

Deletes a topology link

### Syntax

```
openstack network efa-topology-link-map delete [--help ]
```

```
openstack network efa-topology-link-map delete [-h] <link> [<link> ...]
```

Specify the UUID of the link to delete.

### Examples

```
openstack network efa-topology-link-map delete  
08bb90b6-ac37-4b7f-aa80-d1d08de7e54b  
012b90b6-ac12-427f-a220-d1d08de7e123
```

## openstack network efa-topology-link-map list

Lists topology links

### Syntax

```
openstack network efa-topology-link-map list [--help ]
```

### Examples

The following example configures....

```
# openstack network efa-topology-link-map list
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| ID                                     | Host           | Nic | Provider Network |
Switch      | Port | Po Name |                                     |
+-----+-----+-----+-----+-----+-----+
| 08fa6b52-8c25-4193-8e4b-00ce3d81f392 | DCGW1          |     | external_nw      |
10.20.246.8 | 0/14 | lag_4   |                                     |
| 0977490d-f07a-4b2c-b683-8359bde19fcc | niyer-devstack | eth3 | physnet3         |
10.20.246.8 | 0/15 |         |                                     |
| 14486fd7-1bd7-4ba0-a159-b2e76f775fdd | niyer-devstack | eth4 | default          |
10.20.246.8 | 0/16 |         |                                     |
| 4ce69af9-03ce-49d7-a0dc-0b5c64e25ded | niyer-devstack | eth1 | physnet11       |
10.20.246.8 | 0/13 | lag_1   |                                     |
| c3879465-ecdd-492b-918e-981c104005ba | DCGW1          |     | external_nw      |
10.20.246.7 | 0/14 | lag_4   |                                     |
| c9663e8d-f91c-44ff-b4a1-fb97e603bda2 | niyer-devstack | eth0 | physnet11       |
10.20.246.7 | 0/13 | lag_1   |                                     |
| daf3c58a-9b24-4069-afe8-94d2a20004c9 | niyer-devstack | eth2 | physnet2         |
10.20.246.7 | 0/15 |         |                                     |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
```

## openstack router create

---

Creates a router in centralized mode.

### Syntax

```
openstack router create [ router-name --centralized ]
```

### Parameters

*router-name* **--centralized**

Creates a router in centralized mode.

### Usage Guidelines

In centralized mode, routing is configured only on the border leaf pairs.

When you create a router, the default mode is centralized. You can create a router with or without the **--centralized** option and achieve the same result.



#### Note

EFA does not support routers configured in distributed mode, although the feature was supported in a previous release. The functionality to configure distributed mode is turned off.

Each of these examples creates 1 centralized router called R2.

```
$ openstack router create R2  
  
$ openstack router create R2 --centralized
```





# Neutron REST Endpoints

[Neutron REST Endpoints](#) on page 89

This section describes Neutron REST endpoints.

## Neutron REST Endpoints

The following section shows API handled Extreme M12 drivers. For more information on OpenStack APIs, refer to [OpenStack Networking API Guide](#).

### EFA Topology Neutron Extension

```
GET
/v2.0/efa_topologies
Shows details of the topology

Normal response codes: 200
```

**Table 17: Response parameters**

Name	IN	Type	Description
id	Body	UUID	Topology ID
host	Body	String	Host name of the compute
nic	Body	String	NIC on the compute
provider_network	Body	String	Provide network to which the NIC belongs
switch	Body	String	IP address of the switch to which the NIC is connected
Port	Body	String	Switch interface to which the NIC is connected
po_name	Body	String	LAG as created on EFA for the NICs

```
DELETE
/v2.0/efa_topologies{id}
```

Deletes a topology link from its associated resources

**Table 18: Request parameters**

Name	IN	Type	Description
id	Path	UUID	Topology ID

POST

/v2.0/efa\_topologies

Create a link on the efa\_topology object

**Table 19: Request parameters**

Name	IN	Type	Description
host	Body	String	Host name of the compute
nic	Body	String	NIC on the compute
provider_network	Body	String	Provide network to which the NIC belongs
switch	Body	String	IP address of the switch to which the NIC is connected
Port	Body	String	Switch interface to which the NIC is connected
po_name	Body	String	LAG as created on EFA for the NICs

```
body = {'efa_topology':
  { 'host': 'Openstack115',
    'nic': 'eth0',
    'switch': "10.24.35.225",
    'provider_network': "PFPT_LAG",
    'port': '0/1',
    'po_name': 'lag_1',
  }}

```

## EFA Border-Leaf Pair Mapping

GET

/v2.0/efa\_bl\_pairs

Shows details of the EFA border leaf pair  
 Normal response codes: 200

**Table 20: Response parameters**

Name	IN	Type	Description
id	Body	UUID	Border-leaf pair ID
name	Body	String	Name of the border-leaf pair
Switch_ips	Body	String	Comma separated border-leaf ipaddress

DELETE  
 /v2.0/efa\_bl\_pairs{id}  
 Deletes border-leaf pair by ID

**Table 21: Request parameters**

Name	IN	Type	Description
id	Path	UUID	Border-leaf pair ID

POST  
 /v2.0/efa\_bl\_pairs  
 Creates a border-leaf pair object which will be used for L3 configurations

**Table 22: Request parameters**

Name	IN	Type	Description
name	Body	String	Name of the border-leaf pair
Switch_ips	Body	String	Comma separated border-leaf ipaddress

```
body = {'efa_bl_pair': {
    'name': 'BLPAIR1',
    'switch_ips': "10.24.35.225,10.24.35.224",
}}
```



# SR-IOV and Multi-Segment Support

---

[SR-IOV Network](#) on page 92

[Multi-Segment Network](#) on page 95

This section describes SR-IOV and multi-segment support.

## SR-IOV Network

---


Single Root I/O Virtualization (SR-IOV) allows a single physical NIC to appear as multiple physical NICs. All NICs in the VM must be tagged to the appropriate VLAN.

The two SR-IOV functions are as follows:

- Physical Function (PF) - Physical Ethernet controller that supports SR-IOV
- Virtual Function (VF) - Virtual device created from a physical Ethernet controller

The following figure shows a network with:

- ens3f1 (Compute-1) and ens3f0 (Compute-2) as part of network for PF-PT
- ens3f0 (Compute-1) and ens2f0 (Compute-2) as part of network for VF-PT



**Figure 14: SR-IOV network (VF passthrough)**

## Create PCI Passthrough Allowed List

### Procedure

1. Configure the PCI passthrough allowed list in the `/etc/nova/nova.conf` and `/etc/nova/nova-cpu.conf` files on Compute node 1.

```
[default]
pci_passthrough_whitelist = [{"devname": "ens3f1", "physical_network":
"sriovnet2"}, {"devname": "ens3f0", "physical_network":
"sriovnet1"}]
```

2. Restart the Nova server.
3. Repeat the procedure for Compute node 2.

## Configure SR-IOV Agent

### Procedure

1. Configure the SR-IOV NIC Agent in the `/etc/neutron/plugins/ml2/sriov_agent.ini` file on Compute node 1.

```
[securitygroup]
firewall_driver = neutron.agent.firewall.NoopFirewallDriver
[sriov_nic]
physical_device_mappings = sriovnet1:ens3f0,sriovnet2:ens3f1
exclude_devices =
```

2. Run the SR-IOV NIC Agent on Compute node 1.

```
/usr/local/bin/neutron-sriov-nic-agent --config-file /etc/neutron/neutron.conf --
config-file /etc/neutron/plugins/ml2/sriov_agent.ini
```

3. Repeat the procedure for Compute node 2.

## Configure Nova Scheduler

### Procedure

1. Configure the Nova Scheduler in the `/etc/nova/nova.conf` file on both controller and compute nodes.

```
enabled_filters = ...,PciPassthroughFilter
available_filters = nova.scheduler.filters.all_filters
```

2. Restart the Nova Scheduler.

## Configure Mechanism Drivers for SR-IOV

### Procedure

1. Configure the `sriovnicswitch` mechanism driver in the `m12_conf.ini` file on each controller.

```
[m12]
tenant_network_types = vlan
type_drivers = vlan
mechanism_drivers = openvswitch, sriovnicswitch, extreme_efa
[m12_type_vlan]
network_vlan_ranges = physnet1:2:500
```

2. Ensure that `sriovnet` is configured for the selected network type in the `m12_conf.ini` file on each controller.

```
[m12_type_vlan]
network_vlan_ranges = sriovnet1:100:500,sriovnet2:100:500
```

3. Restart the Neutron server.

## Create Network for VF-PT

### Procedure

Create a network for VF-PT.

```
# openstack network create --provider-physical-network sriovnet1 --provider-network-type
vlan --provider-segment 101 \ sriov-net

# openstack subnet create sriov-subnet --network sriov-net --subnet-range 10.65.217.0/24
--gateway 10.65.217.254
```

## Create Network for PF-PT

### Procedure

Create a network for PF-PT.

```
# openstack network create --provider-physical-network sriovnet2 --provider-network-type
vlan --provider-segment 102 \ pt-net

# openstack subnet create pt-subnet --network pt-net --subnet-range 10.65.217.0/24 --
gateway 10.65.217.254 pt_net_id=$(openstack network show pt-net -c id -f value)
```

All corresponding endpoints are provisioned.

## Create Virtual Machines

### Procedure

1. Create a virtual machine on Compute node 1.

```
# openstack server create --flavor ds512M --image vlan-capable-image --nic port-id=
$port_id --availability-zone nova:Compute-1 sriov-instance-1
```

2. Repeat the procedure for Compute node 2.

## Create SR-IOV Direct Ports

### Procedure

1. Create the SR-IOV direct port.

```
net_id=$(openstack network show sriov-net -c id -f value)

# openstack port create --network $net_id --vnic-type direct \ sriov-port port_id=$
(openstack port show sriov-port -c id -f value)
```

2. Repeat the procedure for the second port.

## Create SR-IOV Direct-Physical Port

### Procedure

1. Create a SR-IOV Direct-Physical port.

```
# openstack port create --network $pt_net_id --vnic-type direct-physical pt-port pt_id=
$ (openstack port show pt-port -c id -f value)
```

2. Repeat the procedure to create the second SR-IOV Direct-Physical port.

## Delete SR-IOV Entities

### Procedure

Delete the SR-IOV entities.

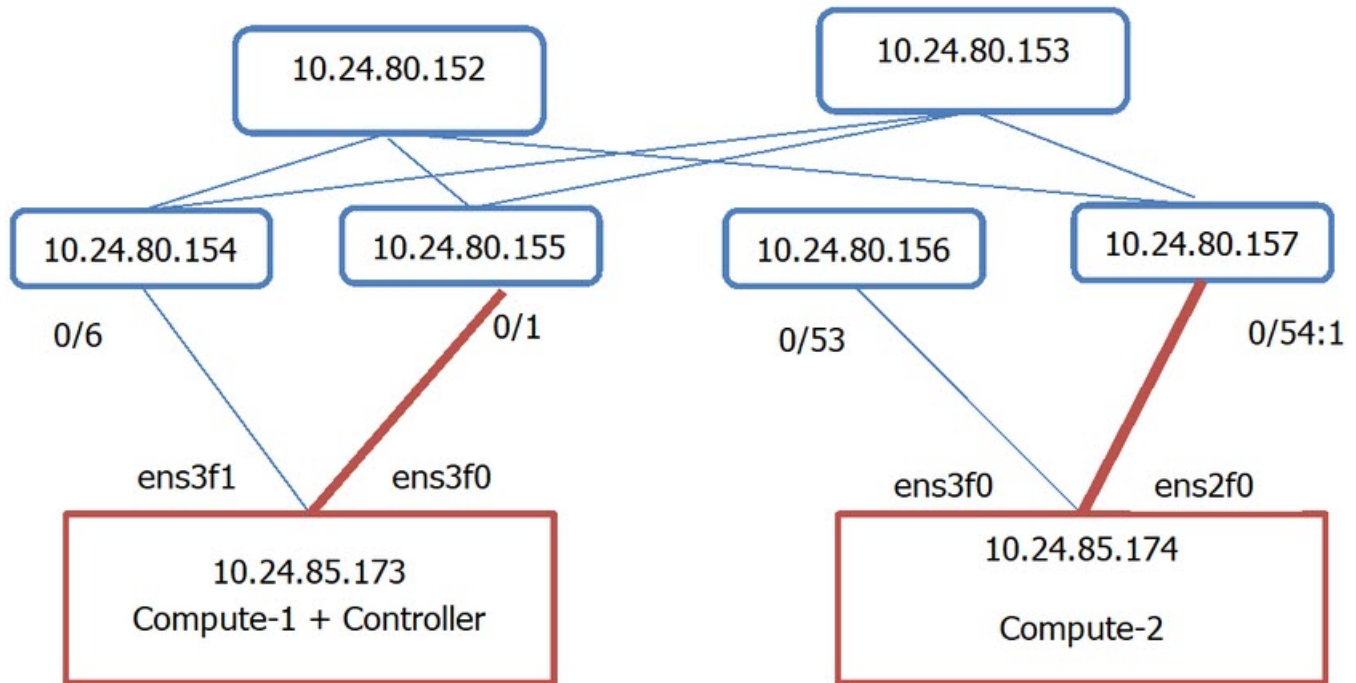
```
# openstack server delete sriov-instance-1
# openstack server delete sriov-instance-2
# openstack port delete sriov-port2
# openstack port delete sriov-port
# openstack subnet delete sriov-subnet
# openstack network delete sriov-net
```

## Multi-Segment Network

A network segment is an isolated Layer 2 segment within a network. A network can contain multiple network segments.

The following figure shows a network with:

- ens3f1 (Compute-1) and ens3f0 (Compute-2) configured as Virtio ports in physnet1
- ens3f0 (Compute-1) and ens2f0 (Compute-2) configured as SRIOV ports in sriovnet1



**Figure 15: Overview of multi-segment network**

## Configure Segments in Neutron

### Procedure

1. Configure segments in the `/etc/neutron/neutron.conf` file.

Ensure that the placement IP address points to the Nova server on the controller node.

```
[DEFAULT]
# ...
service_plugins = ..., segments

[placement]
auth_url = http://10.24.85.173/identity
project_domain_name = Default
project_name = service
user_domain_name = Default
password = apassword
username = nova
auth_url = http://10.24.85.173/identity_admin
auth_type = password
region_name = RegionOne
```

2. Restart the Neutron server.

## Configure Multi-Segment Network

### Procedure

Configure the segment network.

```
# openstack network create --share --provider-physical-network sriovnet2 --provider-
network-type vlan --provider-segment 2016 multisegment
```



## Create Multi-Segment Network

### Procedure

1. Create a multi-segment network.

```
# openstack network create --provider-network-type vlan --provider-segment 333 msnet1
```

2. (Optional) Create a segment for external connectivity through DC Gateway (DC GW).

```
# openstack network segment create --physical-network EXT_A --segment 333 --network-type vlan --network msnet1 extseg
```

Regular Virtio ports can be created on a separate segment of the same network.

All corresponding endpoints are provisioned for SR-IOV PF-PT and DC GW after segment creation.

## Rename Multi-Segment Network

The network must be formed with SR-IOV provider network.

### Procedure

Rename the multi-segment network.

```
# segment1=openstack network segment list -c ID -f value openstack network segment set --name segment1 $virtio_segment
```

## Create Subnet on Segment

### Procedure

Create a subnet on the segment.

```
# openstack subnet create --network multisegment1 --network-segment dhcp_segment --ip-version 4 --subnet-range 203.0.113.0/24 \ multisegment-segment1-v4
```

## Create a Port on SR-IOV Segment

### Procedure

Create a port on SR-IOV segment.

```
# openstack port create --network multisegment --vnic-type direct-physical sriov_port
```

## Create a VM Using the Port

### Procedure

Create a virtual machine using the port.

```
# openstack server create --flavor ds512M --availability-zone nova:Compute-1 --image xenial-server-amd64 --nic port-id=sriov_port sriov_vm
```

```
# ip link add link ens4 name ens4.2016 type vlan id 2016 sudo dhclient ens4.2016
```

```
# sudo dhclient ens4.2016
```