ExtremeCloud Orchestrator CLI Administration Guide

3.2.1

9037725-01 RevAA
May 2023

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

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/

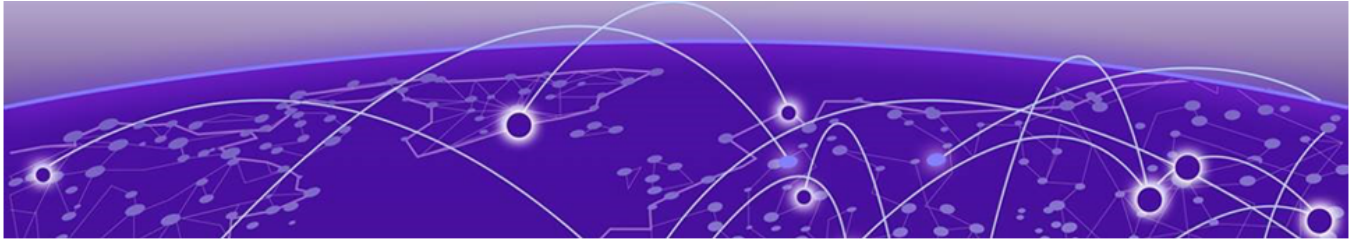

Table of Contents

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 *enter* and *type*	When you see the word *enter* 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 *type*.
Key names	Key names are written in boldface, for example **Ctrl** or **Esc**. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press **Ctrl**+**Alt**+**Del**
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.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
italic text	Italic text indicates variable content.
`[ ]`	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
`{ x | y | z }`	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
`x | y`	A vertical bar separates mutually exclusive elements.
`< >`	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*`[`*member*`...]`.
\	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/.

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

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.

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.

What's New in this Document

The following table describes the information added to this guide for the ExtremeCloud Orchestrator (XCO) 3.2.1 software release.

Table 4: Summary of changes

Feature	Description	Link
Next Hop Recursion	Describes procedure to enable next hop recursion on tenant VRF.	Configure Next Hop Recursion on page 223
BGP Peer Group Activate	Describes procedure to activate BGP peer group.	Activate Peer Group on Tenant BGP on page 341

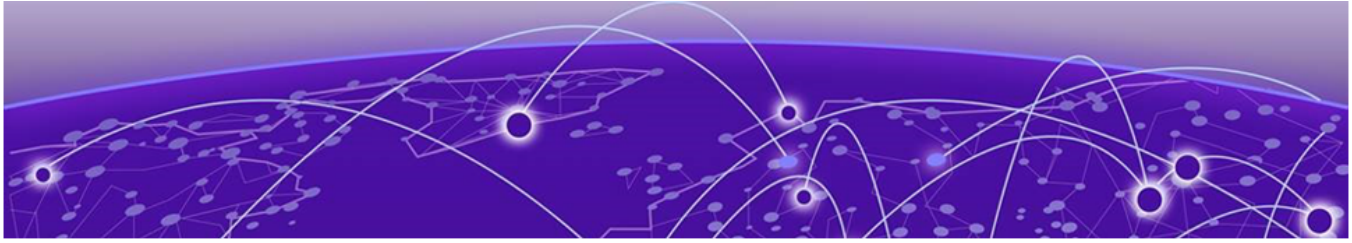

Introduction to ExtremeCloud Orchestrator

ExtremeCloud™ Orchestrator (XCO) is an orchestration application that provides a unified and holistic graphical user interface (GUI) and application programming interface (APIs) for visibility management (visibility skill) and fabric-wide life cycle management (fabric skill) with highly scalable and flexible deployment model for Extreme solutions.

XCO integrates Extreme Fabric Automation (EFA) and Extreme Visibility Manager (XVM) solutions into a single orchestration solution. XCO provides common infrastructure and consistent installation and upgrade strategies for MLX, SLX, Extreme 8000 series, and 9920 devices with focus on scalability and performance.

XCO provides an industry leading user interface with a comprehensive, microservices-based solution to tailor the network to the changing user behavior. The user interface enables IP fabric life-cycle management of SLX, and Extreme 8000 series devices and visibility and policy management of MLX, SLX, and 9920 devices.

Note
All procedures in this document are performed using the CLI commands.

Evolution of EFA and XVM into XCO

This section provides an overview of the evolution of Extreme Fabric Automation (EFA) and Extreme Visibility Manager (XVM) into ExtremeCloud Orchestrator (XCO). This information is intended to help existing EFA and XVM users understand the transformation of EFA and XVM into XCO.

ExtremeCloud™ Orchestrator is a comprehensive microservice based, cloud-native solution that provides organizations the ability to visualize at a workspace level using the user interface or shift to the orchestration level to integrate and automate the network infrastructure through APIs.

XCO integrates Extreme Fabric Automation (EFA) and Extreme Visibility Manager (XVM) solutions:

- EFA: Automates life-cycle management that includes design, deployment, operation, and refresh of IP fabric networks. For more information, see Fabric Automation and Orchestration on page 17.
- XVM: Manages network packet broker (NPB) and visibility solution. For more information, see Visibility Solution on page 18.

XCO Architecture

The XCO architecture is built on the concept of **composable skills** that provide specific capabilities and functions.

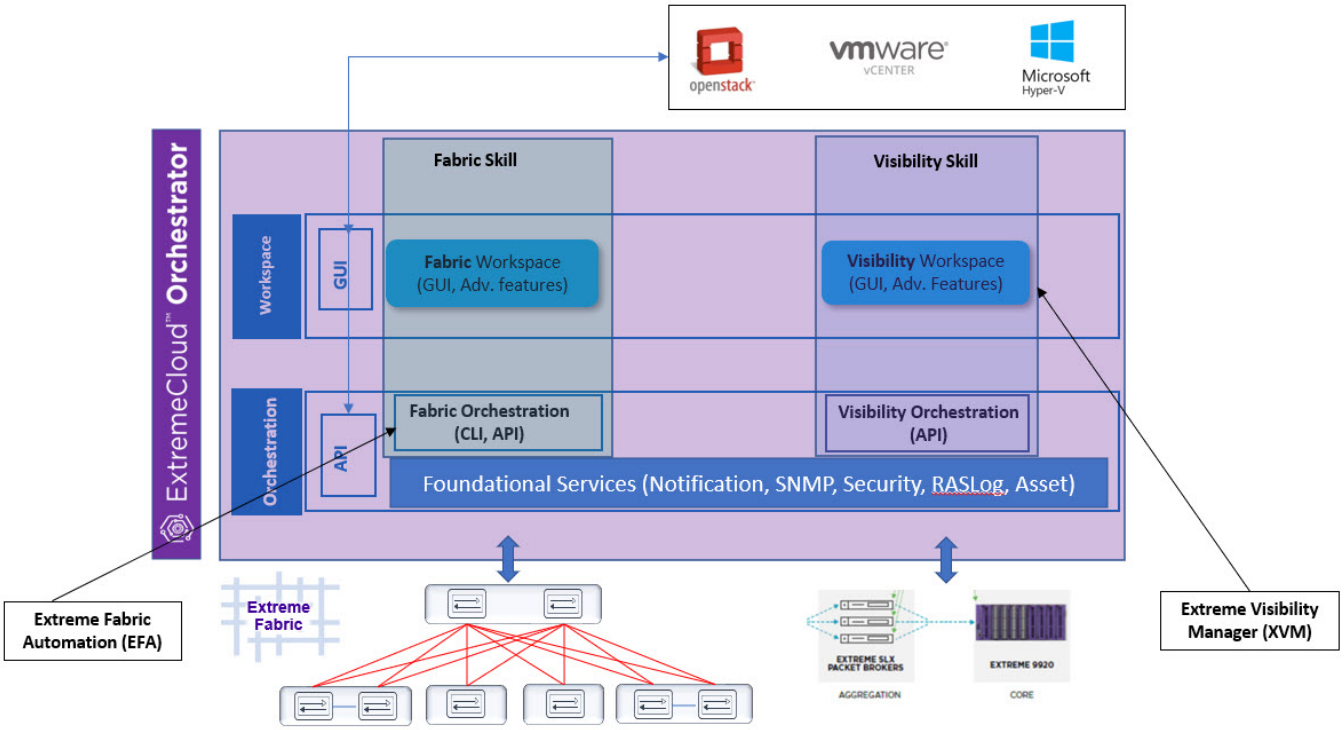

Figure 1: Overview of ExtremeCloud Orchestrator

Skills

The skills within XCO are designed to work together to deliver a unified and integrated orchestration experience for the network. With this approach, organizations can quickly deploy a specific skill and can scale and evolve the orchestration capabilities as needs change over time.

Activation and access of the required skills within XCO can be done at the orchestration level or workspace level:

- Fabric Skill (formerly EFA): Provides the foundation for orchestrating the network to access the complete network as a fabric-like abstraction layer.

- Visibility Skill (formerly XVM): Provides centralized management for all Extreme Networks visibility products.

Retention of EFA CLI and EFA Terminology

The EFA CLI is retained in XCO to ensure backward compatibility for existing EFA users. This allows organizations to continue using the EFA CLI for the life-cycle management of IP fabric, ensuring a smooth transition to XCO. Additionally, the term EFA refers to the fabric skill within XCO.

Fabric Automation and Orchestration

XCO automates and orchestrates SLX IP fabrics, Extreme 8000 series, and tenant networks, with support for the following:

- Building and managing small data center (non-Clos) fabrics and 3-stage and 5-stage IP Clos fabrics
- Managing tenant-aware Layer 2 and Layer 3 networks
- Integrating Virtual Management ecosystem platforms, such as VMWare vCenter, OpenStack, and Microsoft SCVMM
- Providing a single point of configuration for your entire fabric

XCO 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 XCO Microservices on page 18.

CLI and API

Using the built-in command and OpenAPI-based REST APIs, you can discover physical and logical assets, build and manage fabrics, manage the XCO system, and configure security. For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1* and REST API Documentation for XCO on page 23.

Deployment

For more information about deployment scenarios, see the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

XCO on TPVM

TPVM (Third-Party Virtual Machine) is a guest VM that resides on Extreme SLX devices. You can run XCO from the SLX 9150, SLX 9250, Extreme 8520, Extreme 8720, or SLX 9740 TPVM. In this context, XCO leverages the K3S Kubernetes cluster as an underlying infrastructure for the XCO services deployment. The K3S cluster is a single instance and an important component for supporting high availability. A maximum of 24 devices is supported, either 24 devices in one fabric or 24 devices across multiple fabrics.

XCO on an external VM

You can deploy XCO on an external Virtual Machine to support more than 24 devices or based on where tools are deployed in the data center.

XCO for high availability

A high-availability cluster is a group of servers that provide continuous up time, or at least minimum down time, for the applications on the servers in the group. If an application on one server fails, another server in the cluster maintains the availability of the service or application. You can install XCO on a two-node cluster, including on TPVM, for high availability.

Visibility Solution

ExtremeCloud Orchestrator supports several network packet broker devices as part of the visibility solution to provide centralized device and policy management.

For more information, see .

XCO Microservices

XCO 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.

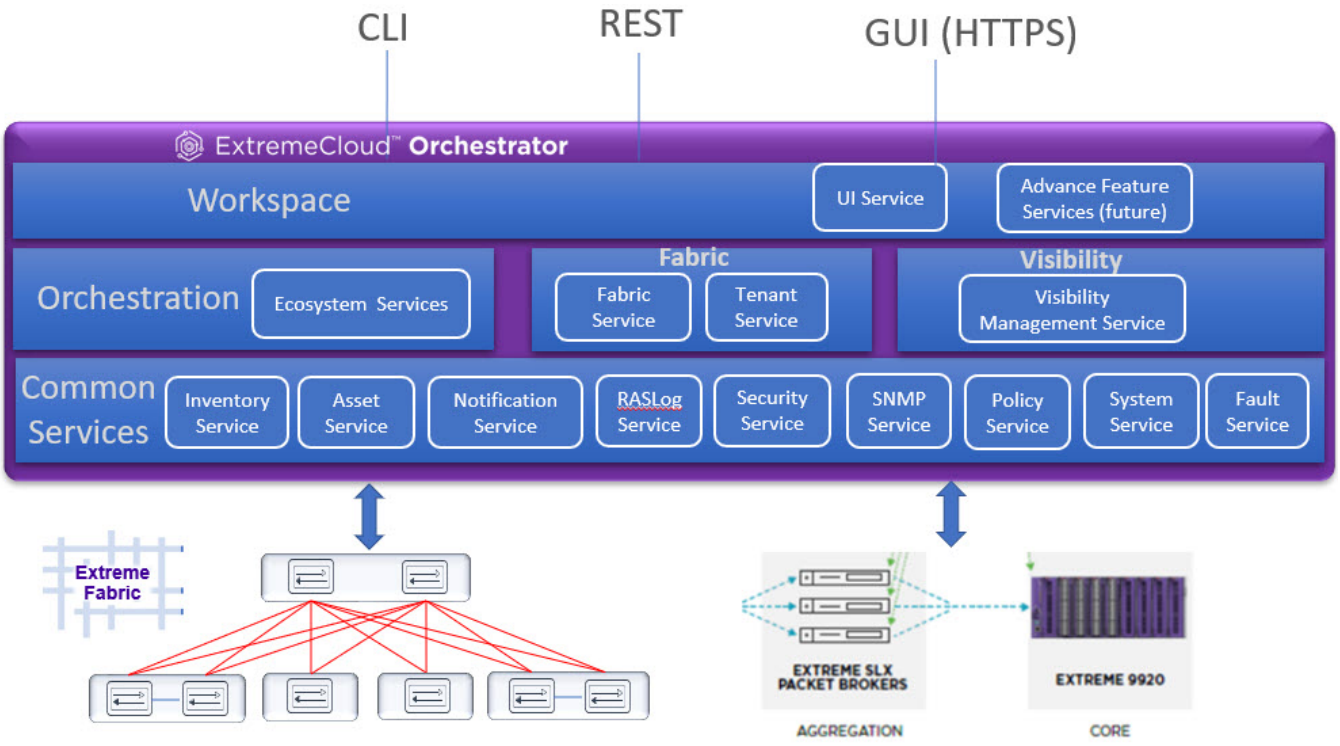

Figure 2: Microservices in the XCO architecture

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 it 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 XCO 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 XCO for SLX-OS commands and push the same configuration to multiple devices.

Asset Service

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, alarms, and tasks to external entities:

- Events: Device events derived from the syslog events received from the managed devices.
- Alerts: Notifications that services in XCO send for unexpected conditions.
- Alarms: A stateful entity that is raised and cleared by the system.
- Tasks: 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 *ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1*.

Security Service

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

- Authentication Service: The service validates user credentials and supports host user login, local user management and remote (LDAP, TACACS) login.
- Authorization Service: The service provides role management and validates the permissions that the user can perform on XCO.

SNMP Service

The SNMP Service processes SNMP traps from devices and forwards notifications to subscribers. For more information, see XCO as SNMP Proxy in the *ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1*.

Policy Service

Policy Service in XCO manages and configures IP prefix lists and route maps on fabric devices. It subscribes to the inventory service to receive events including device registration, device deletion, and changes to previously identified IP prefix lists and route maps.

System Service

The system service provides options to configure system-level settings, such as supportsave, backup, and feature enablement. It periodically takes a backup of the XCO system.

Fault Service

The Fault Service raises alerts and alarms when issues are detected to enable system administrators to monitor and troubleshoot.

Extreme Visibility Manager

Extreme Visibility Manager (Visibility Manager), a Kubernetes-based microservices application, provides centralized device and policy management as part of the Extreme Visibility solution.

Visibility Manager supports several network packet broker devices. Although devices have different functionality and different configuration methods, Visibility Manager seamlessly interacts with all supported devices for simplified management.

You use Visibility Manager to perform much of the same traffic configuration that you might otherwise perform from the command-line interface of your network packet broker operating system. And then you use Visibility Manager to analyze the traffic for insight into issues such as network usage, load-balancing irregularities, and security threats.

Visibility Manager managed objects work together to accomplish most packet broker functions. You configure the objects from the user interface.

For more information, see Extreme Visibility Manager Administration and User Guide Version 6.1.0.

Ecosystem Services

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

VMware vCenter Service

The vCenter integration provides connectivity between XCO and vCenter using a REST API. XCO does not connect to individual ESXi servers. All integration is done through vCenter. For more information, see the . Integration support includes the following:

- Registration or deregistration of one or more vCenter servers in XCO
- Updates for vCenter asset details
- Lists of information about vCenter servers
- Inventory integration
- Dynamic updates about Tenant Service integration from vCenter and from XCO 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 . 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 XCO foundation services in an IP fabric. For more information, see the . 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 provider:physical_network (physnet) value.

    A physnet is the value of a network's provider `:physical_network` attribute.
  - 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 XCO 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

REST API Documentation for XCO

When XCO is installed, REST API documentation is available as an HTML reference: `http://<host_ip>/docs`.

The REST API is specified by OpenAPI and Swagger.

Specific API guides for the XCO services are available on the Extreme Networks website. Select **ExtremeCloud Orchestrator** here: https://www.extremenetworks.com/support/documentation/product-type/software/. And then select the version of XCO you want to work with.

API guides are available for the following services:

- Authentication service
- Authorization service
- Fabric service
- FaultManager service
- Hyper-V service
- Inventory service
- Monitoring service
- Notification service
- OpenStack service
- RASlog service
- RBAC service
- SNMP service
- System service
- Tenant service
- vCenter service
- Policy service

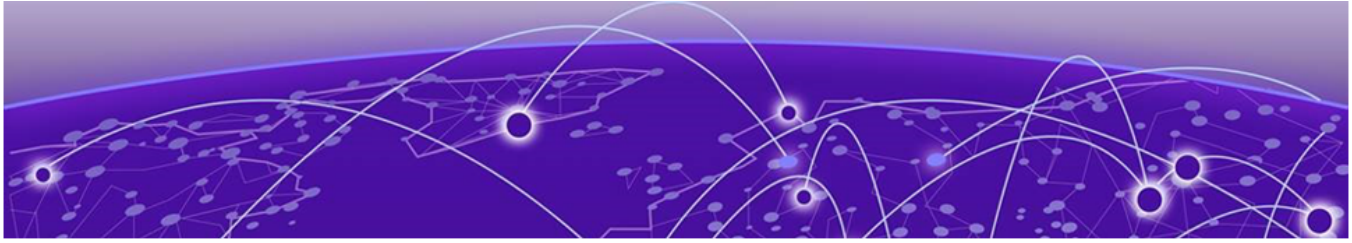

XCO System Management

Verify the Running System and Services

You can use various commands and scripts to verify the status of the XCO system, to help troubleshoot, and to view details of XCO nodes, PODs, and services.

Before You Begin

After any of the following scenarios, wait 10 minutes for XCO micro-services to be operational before you run XCO commands.

- Powering on the OVA
- Rebooting the OVA
- Rebooting the TPVM
- Rebooting the SLX (which also reboots the TPVM)
- Rebooting the server on which the XCO is installed

About This Task

Follow this procedure to verify the status of the XCO system and services.

Procedure

1. Verify the K3s installation in a TPVM.
   a. Run the **`show efa status`** command from the SLX command prompt.

      Ensure that the status of all the nodes are up.

```
device# show efa status
================================================
               EFA version details
================================================
Version : 3.2.0
Build: GA
Time Stamp: 23-03-16:23:17:04
Mode: Secure
Deployment Type: multi-node
Deployment Platform: TPVM
Deployment Suite: Fabric Automation
Deployment IP Mode: ipv4
Virtual IP: 10.20.54.87
Node IPs: 10.20.54.88,10.20.54.89
Node IPv6s: fc00::5:4288:2fff:febd:bc04,fc00::5:4288:2fff:febd:aa04
--- Time Elapsed: 9.30156ms ---

================================================
               EFA Status
================================================
+-----------+---------+--------+-------------+
| Node Name | Role    | Status | IP          |
+-----------+---------+--------+-------------+
| node-1    | active  | up     | 10.20.54.88 |
+-----------+---------+--------+-------------+
| node-2    | standby | up     | 10.20.54.89 |
+-----------+---------+--------+-------------+
--- Time Elapsed: 19.438967114s ---
```

      Output varies by type of deployment, such as single-node or multi-node, and the services that are installed.
2. View details of XCO nodes, PODs, and services.
   a. Run the **`efa status`** command.

      Ensure that the status of all the nodes are up.

      On a multi-node installation:

```
+-----------+---------+--------+---------------+
| Node Name | Role    | Status | IP            |
+-----------+---------+--------+---------------+
| tpvm2     | active  | up     | 10.20.216.242 |
+-----------+---------+--------+---------------+
| tpvm1     | standby | up     | 10.20.216.241 |
+-----------+---------+--------+---------------+
--- Time Elapsed: 4.277420974s ---
```

      On a single-node installation:

```
+-----------+--------+--------+-------------+
| Node Name | Role   | Status | IP          |
+-----------+--------+--------+-------------+
| efa       | active | up     | 10.21.90.43 |
+-----------+--------+--------+-------------+
--- Time Elapsed: 1.461512261s ---
```

      These examples show only a few of all possible rows of detail.

3. Verify that all PODs are in a running state.
   a. Run the **k3s kubectl get pods -n efa** command.

```
(efa:extreme)extreme@node-1:~$ k3s kubectl get pods -n efa -o wide

NAME                          READY STATUS   RESTARTS AGE   IP            NODE
NOMINATED READINESS

NODE      GATES
efa-api-docs-z84wn            1/1   Running  0        5h3m  10.42.194.72  efa
<none>    <none>
gosystem-service-t4h2b        1/1   Running  0        5h3m  10.42.194.74  efa
<none>    <none>
rabbitmq-vn27v                1/1   Running  0        5h4m  10.42.194.69  efa
<none>    <none>
goinventory-service-vpdj7     1/1   Running  0        5h3m  10.42.194.75  efa
<none>    <none>
goauth-service-g76c4          1/1   Running  0        5h3m  10.42.194.71  efa
<none>    <none>
gorbac-service-jzcnf          1/1   Running  0        5h3m  10.42.194.70  efa
<none>    <none>
gofaultmanager-service-wzwgp 1/1   Running  0        5h3m  10.42.194.73  efa
<none>    <none>
gotenant-service-qmvzb        1/1   Running  0        5h3m  10.42.194.78  efa
<none>    <none>
gonotification-service-h9ms2 1/1   Running  0        5h2m  10.20.54.87   efa
<none>    <none>
goraslog-service-rvjsj        1/1   Running  0        5h3m  10.20.54.87   efa
<none>    <none>
gofabric-service-6c4qs        1/1   Running  0        5h3m  10.42.194.76  efa
<none>    <none>
gopolicy-service-g78bh        1/1   Running  0        5h3m  10.42.194.77  efa
<none>    <none>
gosnmp-service-x86sn          1/1   Running  0        5h1m  10.20.54.87   efa
<none>    <none>
(efa:extreme)extreme@node-1:~$
```

4. Verify the status of the Authentication service.
   a. Run the **systemctl status hostauth.service** script.

```
$ systemctl status hostauth.service
hostauth.service - OS Auth Service
Loaded: loaded (/lib/systemd/system/hostauth.service; enabled; vendor preset:
enabled)
Active: active (running) since Thu 2020-04-23 07:56:20 UTC; 23 h ago
Main PID: 23839 (hostauth)
Tasks: 5
CGroup: /system.slice/hostauth.service
        23839 /apps/bin/hostauth

Apr 23 07:56:20 tpvm2 systemd[1]: Started OS Auth Service
```

5. Restart a service using the **efactl restart-service <service-name>** command.
6. Identify the active node that serves as the database for Kubernetes clusters.
   a. Run the **ip addr show** command from all nodes.
   b. Verify that on one of the Ethernet interfaces, the virtual IP address shows up as the secondary IP address.

Log in to XCO

Use of the XCO command line requires a valid, logged-in user.

About This Task

Follow this procedure to log in to the XCO system and services.

Procedure

1. Verify the status of the XCO deployment using one of the following methods.
   - Run the SLX **`show efa status`** command.
   - Run the XCO **`efactl status`** script (or the **`efa status`** command, as an alternative).
2. Log in to XCO.

   ```
   $ efa login --username <username>
   Password: <password>
   ```

   The <username> variable is optional. If you do not provide a user name, log-in defaults to the current (Unix) user.

   With a successful log-in, the command prompt shows the logged-in user in green text. If the log-in is not successful, the command prompt is displayed in red text.
3. To log out of XCO, run the **`efa logout`** command.

XCO Certificate Management

The following certificates in XCO are automatically generated during installation and registration of devices:

- Device Certificates
- XCO Certificates

Device Certificates

Device certificates are installed and configured during the SLX device registration in XCO.

During the registration of an SLX device in XCO, the following certificates are installed on the device:

1. **OAuth Certificate**: The public certificate for verifying an XCO token is copied to the device. This is the JWT Certificate described in XCO Certificates.
2. **Syslog Certificate**: To push messages to XCO over port 6514.
3. **HTTPS Certificate**: To enable secure communication with the clients.

Along with the certificate installation, the following configuration changes are done on the registered device:

1. HTTP mode is disabled on the device, and HTTPS is enabled.
2. OAuth2 is enabled as the primary mode of authentication. Fallback is set to "local login."

Use the **`efa inventory device list`** command to verify the status of the certificates on the device. If the **Cert/Key Saved** column contains "N," then certificates are not installed.

Syslog CA

Use this topic to learn about the third-party certificates for RASlog service (syslog from SLX).

XCO is shipped with default certificates. These are self-signed and the same certificates are used for listening to the syslog messages received from SLX.

```
$ efa inventory device register --ip=10.x.x.x --username=admin --password=password
+----+------------+-----------+-------+--------------+----------+---------+--------+
| ID | IP Address | Host Name | Model | Chassis Name | Firmware | Status  | Reason |
+----+------------+-----------+-------+--------------+----------+---------+--------+
| 1  | 10.x.x.x   | SLX       | 3012  | SLX9250-32C  | 20.2.3d  | Success |        |
+----+------------+-----------+-------+--------------+----------+---------+--------+
Device Details
--- Time Elapsed: 1m6.570042048s ---
```

The syslog certificate on the device is the default CA that XCO contains. XCO Intermediate CA is pushed to SLX for mutual TLS over 6514 port to receive messages from SLX.

```
SLX# show crypto ca certificates
syslog CA certificate(Server authentication):
SHA1 Fingerprint=A3:E8:F6:CB:46:F6:43:C5:D1:90:1F:A7:C6:58:93:29:77:6F:2F:8E
Subject: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com
Issuer: C=US, ST=CA, L=SJ, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=efa.extremenetworks.com/emailAddress=support@extremenetworks.com
Not Before: Feb 20 22:25:26 2020 GMT
Not After : Feb 17 22:25:26 2030 GMT
```

An enhancement updates RASlog service to use the custom certificates that XCO servers use. The certificate CLI on XCO contains a new parameter, which enables you to upload CA.

```
$ efa certificate server --certificate=my_server_162.pem --key=my_server_162.key --
cacert=ca-chain.pem
Please wait as the certificates are being installed...
Certificates were installed!
--- Time Elapsed: 30.946303683s ---
```

If a third-party certificate is installed on XCO along with CA, syslog CA will be pushed to the device instead of the default XCO Intermediate CA.

```
SLX# show crypto ca certificates
syslog CA certificate(Server authentication):
SHA1 Fingerprint=32:70:EB:91:F4:6D:9C:9F:6E:35:E0:00:20:B8:1A:FF:AF:BA:0D:8A
Subject: C=US, O=xyz, OU=abcd, CN=INTERIM-CN
Issuer: C=US, O=xzy, OU=abcd, CN=ROOT-CN
Not Before: Feb 15 14:56:08 2022 GMT
Not After : Nov 11 14:56:08 2024 GMT
```

If you do not provide any CA certificate, the default certificates of XCO are used. If there are already registered devices, then the syslog certificate is automatically updated on these devices.

Expiry and Alerts

Syslog CA has the same expiry as of XCO Intermediate CA or the third-party CA. Legacy notification is sent to the users in case the certificate is going to expire in 30 days. It supports the following alerts which effects the health of XCO security subsystem.

- DeviceCertificateExpiryNoticeAlert
- DeviceCertificateExpiredAlert
- DeviceCertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

When and intermediate CA is renewed on XCO, it is pushed to SLX.

OAuth Certificate

JWT Verifier from XCO is pushed to SLX during registration.

```
SLX# show crypto ca certificates
oauth2 certificate(OAuth2 token signature validation):
SHA1 Fingerprint=57:55:2F:7A:F0:DB:23:CF:37:67:8D:AE:82:35:D8:2D:18:00:17:9E
Subject: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=extremenetworks.com
Issuer: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=extremenetworks.com
Not Before: Sep 2 13:26:27 2022 GMT
Not After : Aug 30 13:26:27 2032 GMT
```

Expiry and Alerts

Legacy notification is sent to the user if the certificate is going to expire in 30 days. It supports the following alerts which effects the health of XCO security subsystem:

- DeviceCertificateExpiryNoticeAlert
- DeviceCertificateExpiredAlert
- DeviceCertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Upload or Renewal

To upload the token signing certificate to the device, run the following command:

```
(efa:extreme)extreme@tpvm:~$ efa certificate device install --ip=10.x.x.x --certtype=
token
+-------------+---------+
| IP Address  | Status  |
+-------------+---------+
| 10.x.x.x    | Success |
+-------------+---------+
---Time Elapsed: 27.233017418s ---
```

For more information about updating the certificates, see Manual Installation of Certificates on Devices on page 31.

On renewal of certificate, `CertificateRenewalAlert` is raised which changes the health of the system to green.

HTTPS Certificates

When you register a device in XCO, a new certificate is generated for the HTTPS server of SLX device. The certificate is generated with the default CA that XCO contains.

The following is an example of a certificate on SLX after device registration:

```
slx-171# show crypto ca certificates

Certificate Type: https; Trustpoint: none
certificate:
SHA1 Fingerprint=C1:F1:2C:BF:1A:47:7B:46:5D:8F:18:99:0E:58:CF:31:8C:58:5F:CC
Subject: CN=slx-10.x.x.x.extremenetworks.com
Issuer: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com
Not Before: Jan 10 11:12:18 2022 GMT
Not After : Jan 10 11:12:18 2024 GMT
```

You can use the CLI command only to install third-party certificates on a single device at once.

```
(efa:extreme)extreme@tpvm:/apps/test/certs$ efa certificate device install --ip=10.x.x.x
        --cert-type https --https-certificate server.crt --https-key
      my_server.key

WARNING: This will restart the HTTP service on the devices and services will not be able
to connect till the operation is complete. Do you want to proceed [y/n]?

y
+--------------+---------+
| IP Address   | Status  |
| 10.20.61.171 | Success |
+--------------+---------+
--- Time Elapsed: 38.516844258s ---
```

The device must have the new certificates uploaded:

```
slx-171# show crypto ca certificates

Certificate Type: https; Trustpoint: none
certificate:
SHA1 Fingerprint=D8:49:5F:12:AC:FE:BB:CB:95:C2:AC:6B:AF:B6:5B:9E:24:66:59:7D
Subject: CN=10.x.x.x/subjectAltName=IP=10.20.61.171
Issuer: C=US, O=xyz, OU=abcd, CN=INTERIM-CN
Not Before: Feb 10 11:23:36 2022 GMT
Not After : Jun 25 11:23:36 2023 GMT
```

Expiry and Alerts

The HTTPS certificate generated for SLX has an expiry of two years from the date of registration. The device shows the following error message when an HTTP certificate expires:

```
1022 AUDIT, 2025/06/24-17:20:52 (GMT), [SEC-3112], INFO, SECURITY, admin/admin/127.0.0.1/
http/REST Interface,, SLX, Event: X509v3, Certificate Validation failed, Info: Reason =
certificate has expired,
Certificate Details = [Subject CN efa.extremenetworks.com,
Serial 16193545342960822577 Issuer /C=US/ST=CA/O=Extreme Networks/OU=Extreme Fabric
Automation Intermediate/CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com].
```

Legacy notification is sent to the users if the certificate is going to expire in 30 days. It supports the following alerts which effects the health of XCO security subsystem:

```
-    DeviceCertificateExpiryNoticeAlert
-    DeviceCertificateExpiredAlert
-    DeviceCertificateUnreadableAlert
```

For more information, see Fault Management - Alerts on page 506.

Upload or Renewal

To upload the HTTPS certificate to the device, use the following command:

```
(efa:extreme)extreme@tpvm:~$ efa certificate device install --ip=10.x.x.x --certtype=
https

WARNING: This will restart the HTTP service on the devices and services will not be able
to connect till the operation is complete. Do you want to proceed [y/n]?
y
+-------------+---------+
| IP Address  | Status  |
+-------------+---------+
| 10.x.x.x    | Success |
+-------------+---------+
---Time Elapsed: 27.233017418s ---
```

For more information about updating the certificates, see Manual Installation of Certificates on Devices on page 31.

On renewal of certificate, `CertificateRenewalAlert` is raised which changes the health of the system to green.

Manual Installation of Certificates on Devices

You can upload HTTPS and Token certificate on the devices using the following command:

```
efa certificate device install --help
Install certificates on devices

Usage:
  efa certificate device install [flags]

Flags:
      --ip string                 Comma separated range of device IP addresses.
                                  Example: 1.1.1.1-3,1.1.1.2,2.2.2.2
      --fabric string             Specify the name of the fabric
      --cert-type string          Certificate Type (https | token)
      --https-certificate string  Local path to the certificate pem file
      --https-key string          Local path to the key pem file
      --force                     Update the certificate even if already present
  --- Time Elapsed: 3.350424ms ---
```

Note

Fabric and multiple IP do not work with `https|token` (**`efa certificates device install --ips <ip-adddr> certType [ http|token]`**).

Use the following command to install the certificates on multiple devices:

```
efa certificates device install --ip 10.139.44.147-148 --certType https
+---------------+---------+
| IP Address    | Status  |
+---------------+---------+
| 10.139.44.148 | Success |
+---------------+---------+
| 10.139.44.147 | Success |
+---------------+---------+
```

Use the following command to install the HTTPS certificates on the devices in fabric fabric1. If the force option is used, it will update the certificates even if already present:

```
efa certificates device install --fabric fabric1 --certType https --force
+---------------+---------+
| IP Address    | Status  |
+---------------+---------+
| 10.139.44.148 | Success |
+---------------+---------+
| 10.139.44.147 | Success |
+---------------+---------+
```

When you enter the force option, certificates on the devices of interest are updated whether they currently exist or not. If you do not enter the force option, the update reverts to only installing certificates on input devices that do not have them.

Example:

```
Certificates on SLX 10.139.44.147 before and after force:
SLX# show crypto ca certificates
Certificate Type: https; Trustpoint: none
certificate:
SHA1 Fingerprint=CA:7D:13:C6:44:05:71:24:6B:BC:D4:C2:75:95:B6:53:AE:74:03:C0
Subject: CN=slx-10.139.44.147.extremenetworks.com
Issuer: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com
Not Before: Aug 2 13:42:05 2022 GMT
Not After : Aug 2 13:42:05 2024 GMT
syslog CA certificate(Server authentication):
SHA1 Fingerprint=C4:23:B1:A9:6B:DD:45:6C:AA:9B:85:10:63:65:0E:02:77:7D:68:49
Subject: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com
Issuer: C=US, ST=CA, L=SJ, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=efa.extremenetworks.com/emailAddress=support@extremenetworks.com
Not Before: Sep 2 13:14:01 2022 GMT
Not After : Aug 30 13:14:01 2032 GMT
oauth2 certificate(OAuth2 token signature validation):
SHA1 Fingerprint=57:55:2F:7A:F0:DB:23:CF:37:67:8D:AE:82:35:D8:2D:18:00:17:9E
Subject: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=extremenetworks.com
Issuer: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=extremenetworks.com
Not Before: Sep 2 13:26:27 2022 GMT
Not After : Aug 30 13:26:27 2032 GMT

SLX# show crypto ca certificates
Certificate Type: https; Trustpoint: none
certificate:
SHA1 Fingerprint=73:06:CD:84:F3:C9:12:49:70:88:57:4A:A5:97:43:91:6A:BA:98:A1
Subject: CN=slx-10.139.44.147.extremenetworks.com
Issuer: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com
Not Before: Aug 2 13:44:24 2022 GMT
```

```
Not After : Aug 2 13:44:24 2024 GMT
syslog CA certificate(Server authentication):
SHA1 Fingerprint=C4:23:B1:A9:6B:DD:45:6C:AA:9B:85:10:63:65:0E:02:77:7D:68:49
Subject: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA/emailAddress=support@extremenetworks.com
Issuer: C=US, ST=CA, L=SJ, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=efa.extremenetworks.com/emailAddress=support@extremenetworks.com
Not Before: Sep 2 13:14:01 2022 GMT
Not After : Aug 30 13:14:01 2032 GMT
oauth2 certificate(OAuth2 token signature validation):
SHA1 Fingerprint=57:55:2F:7A:F0:DB:23:CF:37:67:8D:AE:82:35:D8:2D:18:00:17:9E
Subject: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=extremenetworks.com
Issuer: C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation,
CN=extremenetworks.com
Not Before: Sep 2 13:26:27 2022 GMT
Not After : Aug 30 13:26:27 2032 GMT
```

XCO Certificates

Multiple certificates are generated and used across the components in XCO.

1. **App Server Certificate**: The certificate of XCO server for secure communication with the clients. This certificate is used on port 443 (default XCO), 8078 (monitor service of XCO), and 6514 (syslog listener on XCO).
2. **Intermediate CA Certificate**: Certificate Authority, which is the issuer of client and server certificates of XCO and HTTPS certificate of SLX.
3. **Root CA Certificate**: Certificate Authority, which is the issuer of Intermediate CA certificate.
4. **JWT Certificate**: The RSA public key for JWT verification. This is also used to send user context from XCO to SLX.
5. **K3s Server Certificate (Internal)**: XCO uses K3s for management of services. This certificate is for secure communication of k3s with clients
6. **K3s CA Certificate (Internal)**: XCO uses K3s for management of services. These certificates are used for generating all the certificates of K3s.
7. **Host Authentication Service Certificate (Internal)**: The server certificate of host authentication service on XCO

The following tables provide information about XCO certificates.

For Alerts related to Alarm/Notifications, refer to the alerts described in Fault Management - Alerts on page 506.

Certificate	Location in TPVM deployment	Location in server deployment	Description	Default Validity Period	Impact on the system	Renewal Procedure	Alarm/ Notification
SSL/TLS Certificate of XCO	/apps/ efadata/ certs/own/ tls.crt	/opt/ efadata/ certs/own/ tls.crt	The certificate of XCO server for secure communication with the clients. The same certificate is used on port 443 (default XCO), 8078 (monitor service of XCO), 6514 (syslog listener on XCO), 8079 (host authentication service of XCO)	Expires in 3 years from installation. Reset after every subinterface creation/ upgrade	If the certificate expires, then the server communication with SSL verification enabled will fail. Disables syslog messages from the devices	Use the **`efa certificate server renew`** command as described in the XCO Server Certificate on page 38.	Notification is sent to XCO subscribers from 30 days to expiry and warning message on every login from 7 days to expiry. Notification is sent to XCO subscribers: 1. After 30 days of expiry 2. Expired certs 3. Renewal certs
K3s CA Certificate	/apps/ rancher/k3s/server/tls/ server-ca.crt	/var/lib/ rancher/k3s/server/tls/ server-ca.crt	XCO uses K3s for management of services. These certificates are for secure communication of K3s with clients.	Expires in 10 years from the date of installation.		K3s CA on page 42.	Notification is sent to XCO subscribers: 1. After 30 days of expiry 2. Expired certs 3. Renewal certs

Certificate	Location in TPVM deployment	Location in server deployment	Description	Default Validity Period	Impact on the system	Renewal Procedure	Alarm/ Notification
Intermediate CA Certificate of XCO	/apps/efadata/certs/ca/extreme-ca-cert.pem	/opt/efadata/certs/ca/extreme-ca-cert.pem	The certificate of Certificate Authority, which is the issuer of client and server certificates of XCO and HTTPS certificate of SLX. Same certificate is seen as SyslogCA on SLX	Expires in 10 years from the date of installation		XCO Intermediate CA on page 37	Not available Notification is sent to XCO subscribers: 1. After 30 days of expiry 2. Expired certs 3. Renewal certs
Root CA Certificate of XCO	/apps/efadata/certs/ca/extreme-ca-root.pem	/opt/efadata/certs/ca/extreme-ca-root.pem	The certificate of Certificate Authority, which is the issuer of Intermediate CA certificate	Expires in 20 years from the date of installation		XCO Root CA on page 36	XCO Certificate Expiry Notice XCO Certificate Expired XCO Certificate Upload or Renewal
HTTPS Certificate of SLX	/apps/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem	/opt/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem	The certificate of SLX Web Server (Apache) for secure communication with the device from XCO	Expires in 2 years from installation	System will not use encryption for HTTPS requests	HTTPS Certificates on page 30	Notification is sent to XCO subscribers from 30 days of expiry.

Certificate	Location in TPVM deployment	Location in server deployment	Description	Default Validity Period	Impact on the system	Renewal Procedure	Alarm/ Notification
K3s Certificate - XCO internal	/apps/rancher/k3s/server/tls/	/var/lib/rancher/k3s/server/tls/	XCO uses k3s for management of services. This certificate is for secure communication of k3s with clients	Expires in 1 year from installation. Reset after every upgrade of XCO		K3s Server Certificate on page 42	XCO Certificate Expiry Notice
JWT Signing/ Verification - XCO internal	/apps/efadata/certs/cert.crt.pem	/opt/efadata/certs/cert.crt.pem	The RSA public key for JWT verification. This is also used to send user context from XCO to SLX. Same certificate is seen as Oauth certificate on SLX	Expires in 10 years from the date of installation	Disables login to XCO	JWT Certificate on page 41	XCO Certificate Expiry Notice Managed Device Certificate Expiry Notice Managed Device Certificate Expired XCO Certificate Upload or Renewal Managed Device Certificate Upload or Renewal

XCO Root CA

XCO is shipped with Root CA that is used to generate Intermediate CA. The Root CA is unique across each XCO and is generated during installation.

Location

- **TPVM:** `/apps/efadata/certs/ca/extreme-ca-root.pem`
- **Server:** `/opt/efadata/certs/own/extreme-ca-root.pem`

Expiry and Alerts

The XCO Root CA is valid for 20 years from the date of installation. It supports the following alerts which effects the health of XCO security subsystem:

- CertificateExpiryNoticeAlert
- CertificateExpiredAlert
- CertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

You can renew or regenerate the root CA by using either script or command.

To renew or regenerate the Root CA, run the renewal script **`efa_renew_certs.sh`**.

```
sudo bash <path to the script>/efa_renew_certs.sh --type rootca
```

To renew or regenerate the Root CA, run the **`efa certificate server renew`** command.

```
efa certificate server renew --cert-type
```

Note
In TPVM, the renewal script and command are available in the `/apps/efa/` and `/opt/efa/` directory of a server.

After the Root CA is updated,

- New Intermediate CA is generated
- New XCO Server Certificate is generated. If a third-party certificate is used, then the server certificate generation is skipped.

On renewal of certificate, a `CertificateRenewalAlert` is raised which changes the health of the system to green.

XCO Intermediate CA

XCO is shipped with Intermediate CA that is used to

1. Generate server certificate of XCO
2. Generate HTTPS certificate of SLX
3. Connect from Syslog server of SLX

During an upgrade, the old certificates are retained, and will not be regenerated.

Location

- **TPVM:** `/apps/efadata/certs/ca/extreme-ca-cert.pem`
- **Server:** `/opt/efadata/certs/own/extreme-ca-cert.pem`

Expiry and Alerts

The XCO Intermediate CA is valid for 10 years from the date of installation. It supports the following alerts which effects the health of XCO security subsystem:

- CertificateExpiryNoticeAlert
- CertificateExpiredAlert
- CertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

You can renew or regenerate the Intermediate CA by using either script or command.

To renew or regenerate the Intermediate CA, run the renewal script `efa_renew_certs.sh`.

```
sudo bash <path to the script>/efa certificate server renew.sh --type intermediateca
```

To renew or regenerate the Intermediate CA, run the **`efa certificate server renew`** command.

```
efa certificate server renew --cert-type
```

Note
In TPVM, the renewal script and command are available in the `/apps/efa/ and /opt/efa/` directory on a server installation.

After the Intermediate CA certificate is updated,

- A new XCO server certificate is generated. If a third-party certificate is used, then the server certificate generation is skipped.
- The Syslog certificates for the registered devices are automatically updated.
- You must manually update the HTTPS certificate on the devices.

For more information about updating the certificates, see HTTPS Certificates on page 30 for SLX.

On renewal of certificate, `CertificateRenewalAlert` is raised which changes the health of the system to green.

XCO Server Certificate

XCO is shipped with a self-signed certificate that is generated during installation. It is signed by the XCO Intermediate CA on page 37 certificate. This certificate is used on the following ports:

- **443**: Secure installation of XCO
- **8078**: Monitoring service of XCO
- **6514**: RASlog server on port 6514 to connect with devices

Third-party Certificate

You can replace server certificate with a third-party certificate acquired through trusted CAs (for example, Verisign or GoDaddy). The third-party certificate must be present in the host device that is running XCO. You can then install it with the following command:

```
$ efa certificate server --help
Install certificates for EFA

Usage:
  efa certificate server [flags]
  efa certificate server [command]

Available Commands:
  renew       Renew certificates for EFA

Flags:
      --certificate string   Certificate for EFA
      --key string           Key File for the certificate
      --cacert string        CA Certificate File
```

Example:

```
$ efa certificate server --certificate=my_server.pem --key=my_server.key --cacert=ca-
chain.pem
Please wait as the certificates are being installed...
Certificates were installed!
--- Time Elapsed: 30.946303683s ---
```

Note

- If you install your own server certificate to use with the XCO HTTPS server, be sure to reinstall the certificate when you upgrade XCO.
- Generate the third-party certificates and keys without a passphrase. Certificate installation may fail if you generate the third-party certificates and keys with passphrase.
- Ensure that the certificate that is uploaded has validity of at least 90 days.
- XCO relies on common name and the SAN IPs of the certificate. For a single-node deployment, the SAN IP field must have the management IP of the system. In multi-node deployment, ensure that the node IPs and the VIP are present.
- If there are any multiaccess subinterfaces, be sure to add these IPs to the SAN IPs when you generate a certificate.

To upload third-party certificates for HTTPS server on SLX, use the following CLI command. This works only to install certificates on a single device at once.

```
(efa:extreme)extreme@tpvm:/apps/test/certs$ efa certificate device install --ip=10.x.x.x
--cert-type https --https-certificate server.crt --https-key my_server.key

WARNING: This will restart the HTTP service on the devices and services will not be able
to connect till the operation is complete. Do you want to proceed [y/n]?
y
+--------------+---------+
| IP Address   | Status  |
| 10.20.61.171 | Success |
```

```
+--------------+---------+
--- Time Elapsed: 38.516844258s ---
```

The device must have the new certificates uploaded.

```
slx-171# show crypto ca certificates
Certificate Type: https; Trustpoint: none
certificate:
SHA1 Fingerprint=D8:49:5F:12:AC:FE:BB:CB:95:C2:AC:6B:AF:B6:5B:9E:24:66:59:7D
Subject: CN=10.x.x.x/subjectAltName=IP=10.20.61.171
Issuer: C=US, O=xyz, OU=abcd, CN=INTERIM-CN
Not Before: Feb 10 11:23:36 2022 GMT
Not After : Jun 25 11:23:36 2023 GMT
```

XCO utilizes the third-party certificates for northbound access. Prior to XCO 3.2.0, when you run any upgrade or node-replacement procedure, the third-party certificate is replaced with the default certificates of XCO.

It retains the certificates that you have installed during any deployment activities.

In case of any issues while installing the third-party certificates, it will revert back to use the default certificates that are shipped with XCO. The validity of the third-party certificates is verified during XCO upgrade and initial upload of the third-party certificates.

Location

- Default certificate
  - **TPVM:** `/apps/efadata/certs/own/tls.crt`
  - **Server:** `/opt/efadata/certs/own/tls.crt`
- Third-party Certificate
  - **TPVM:** `/apps/efadata/certs/thirdparty/tls.crt`
  - **Server:** `/opt/efadata/certs/thirdparty/tls.crt`
- Third-party CA Certificate
  - **TPVM:** `/apps/efadata/certs/thirdparty/custom-ca-chain.pem`
  - **Server:** `/opt/efadata/certs/thirdparty/ custom-ca-chain.pem`

Expiry and Alerts

The certificate is valid for 3 years from the date of installation. It is regenerated whenever a new multiaccess subinterface is created or deleted from XCO.

Legacy notification is sent to the user if the certificate is going to expire in 30 days. If you do not renew the certificates within 7 days of expiry, a following warning message is displayed on every login to the XCO CLI.

```
(efa:extreme)extreme@tpvm:/apps/test/certs$ efa login
Password:
Login successful.
Warning: The certificate for 'EFA' will expire on '2022-04-08 14:43:43 +0530 IST'.
--- Time Elapsed: 5.532391719s ---
```

XCO server certificate supports the following alerts which effects the health of XCO security subsystem.

- CertificateExpiryNoticeAlert

- CertificateExpiredAlert
- CertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

To renew the server certificate, use the following command:

```
(efa:extreme)extreme@tpvm:/apps$ efa certificate server renew
Certificate renewal is successful
--- Time Elapsed: 33.516064167s ---
```

Note

- Renewal is not applicable if the third-party certificates are installed on the system. You must upload a new certificate as described in the "Third-party certificates" section of HTTPS Certificates on page 30.
- On renewal of certificate or a successful upload, `CertificateRenewalAlert` is raised which changes the health of the system to green.

JWT Certificate

XCO uses JSON Web Tokens for authentication which uses RSA key pair for signing and verification of the tokens.

Location

- **TPVM:** `/apps/efadata/certs/cert.crt.pem`
- **Server:** `/opt/efadata/certs/cert.crt.pem`

Expiry and Alerts

The certificate is valid for 10 years from the date of installation. It supports the following alerts which effects the health of XCO security subsystem:

- CertificateExpiryNoticeAlert
- CertificateExpiredAlert
- CertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

To renew or regenerate token signing certificate, use the following command:

```
(efa:extreme)extreme@tpvm:/apps$ efa certificate server renew --cert-type=token
Certificate renewal is successful.
--- Time Elapsed: 27.233017418s ---
```

After the token certificate is updated, it has to be pushed to all the registered devices. For more information about updating the certificates, see OAuth Certificate on page 29 for SLX.

On renewal of the certificate, `CertificateRenewalAlert` is raised which changes the health of the system to green.

K3s CA

XCO uses K3s for management of microservices which comes up with its own certificates.

Location

- **TPVM:** `/apps/rancher/k3s/server/tls/server-ca.crt`
- **Server:** `/var/lib/rancher/k3s/server/tls/server-ca.crt`

Expiry and Alerts

The certificate is valid for 10 years from the date of installation and is regenerated after every upgrade. It supports the following alerts which effects the health of XCO security subsystem:

- CertificateExpiryNoticeAlert
- CertificateExpiredAlert
- CertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

You can renew or regenerate the K3s CA by using either script or command.

To renew or regenerate the K3S CA, use the renewal script **`efa_k3s_renew_certs.sh`**.

```
sudo bash <path to the script>/efa_k3s_renew_certs.sh --type ca
```

To renew or regenerate the K3S CA, use the **`efa certificate server renew`** command.

```
efa certificate server renew --cert-type
```

Note

- In TPVM, the renewal script and command are available in /apps/efa/ and /opt/efa/ on a server installation.
- If there are any third-party certificates already installed on XCO, reinstall these certificates after K3s CA certificates are regenerated.

On renewal of the certificate, `CertificateRenewalAlert` is raised which changes the health of the system to green.

K3s Server Certificate

XCO uses K3s for management of microservices which comes up with its own certificates.

Location

- **TPVM:** `/apps/efadata/certs/ca/extreme-ca-cert.pem`
- **Server:** `/opt/efadata/certs/own/extreme-ca-cert.pem`

Expiry and Alerts

The certificate is valid for one year from the date of installation which is reset on every upgrade. It supports the following alerts which effects the health of XCO security subsystem:

- CertificateExpiryNoticeAlert
- CertificateExpiredAlert
- CertificateUnreadableAlert

For more information, see Fault Management - Alerts on page 506.

Renewal

You can renew or regenerate the K3s CA by using either script or command.

You can perform the renewal of K3s Server certificate only when:

- K3s server certificate has expired
- K3s server certificates expiry is less than 90 days

Note
In TPVM, the renewal script and command are available in the `/apps/efa/` and `/opt/efa/` directory on a server installation.

To renew or regenerate the K3S server certificate, use the renewal script **`efa_k3s_renew_certs.sh`**.

```
sudo bash <path to the script>/efa_k3s_renew_certs.sh --type server
```

To renew or regenerate the K3S server certificate, use the **`efa certificate server renew`** command.

```
efa certificate server renew --cert-type
```

On renewal of the certificate, `CertificateRenewalAlert` is raised which changes the health of the system to green.

Host Authentication Certificate

Before XCO 3.2.0 release, Host Authentication service runs on the XCO host with its own certificate with an expiry of 10 years. This certificate had no renewal procedure.

The Host Authentication service runs on port 8079. The port is closed from external access. This service reuses the XCO server certificate on 443 with an expiry of 3 years.

All the operations that are performed on the server certificate are applied on the Host Authentication service. This includes uploading of third-party certificate, renewal and so on.

Important
From XCO 3.2.0, the Host Authentication Certificate is not present anymore. You can use the XCO server certificate.

External Certificates

The following tables provide information about external certificates.

Certificate	Location in TPVM deployment	Location in server deployment	Description
Syslog CA (Notification subscriber)	Notification service database in mariadb	Notification service database in mariadb	Connect to external Syslog server for sending notifications over RELP.
Webhook CA	Notification service database in mariadb	Notification service database in mariadb	Connect to external Syslog server for sending notifications over webhooks.
LDAP CA	Auth service database in mariadb	Auth service database in mariadb	Connect to external LDAP server

Certificate Troubleshooting

Issue	Resolution
My device is registered but the certificates do not appear on the SLX device.	Try the following: • Ensure that the device is running at least SLX-OS 20.1.x. • Ensure that the time on the SLX device and the time on the XCO host device are synchronized. • Ensure that the certificates are installed. Run the **efa certificate device install** command.

How do I know about the certificate expiry in XCO?	• Run the following REST API to get the expiry date of all the certificates of XCO:
	`curl -X GET 'https://<vip>:8078/v1/monitor/certificate/expiry' --header 'Authorization:Bearer eyJhbGciOiJSUzI… '.`
	• Run the following openssl command:
	`extreme@tpvm:~$ openssl x509 -in <Location of the certificate> -noout -enddate`
	• Run the **efa certificate expiry show** command.

Issue	Resolution
How do I verify the certificate provided by XCO through its ingress interface?	Run the following command. The output should indicate that `efa.extremenetworks.com` is present. `$ openssl s_client -connect <EFA_IP_ADDR>:443`

There is a security violation on the switch when XCO (installed on TPVM) logs in and tries to access the switch with different usernames. You observe the following logs on SLX console:
`1018 AUDIT, 2021/10/14-17:26:57 (GMT), [SEC-3021], INFO, SECURITY, extreme/root/ 10.20.32.141/ssh/CLI,, SLX, Event: login, Status: failed, Info: Failed login attempt through REMOTE, IP Addr: 10.20.32.141`
`1017 AUDIT, 2021/10/14-17:26:55 (GMT), [SEC-3020], INFO, SECURITY, admin/admin/ 10.20.32.141/ssh/CLI,, SLX8720-32C, Event: login, Status: success, Info: Successful login attempt via REMOTE, IP Addr: 10.20.32.141`
`1002 AUDIT, 2021/10/14-17:26:41 (GMT), [SEC-3020], INFO, SECURITY, admin/admin/ 10.20.32.141/ssh/CLI,, SLX8720-32C, Event: login, Status: success, Info: Successful login attempt via NETCONF, IP Addr: 10.20.32.141`
	Try the following:
	• Ensure that you have correctly followed the system restore process.
	• Ensure that all the devices are registered.
	• Ensure that the certificates are installed on the devices to enable secure connections. Run the **`efa certificate device install --ips <ip-adddr> certType [ http| token]`** command to install the HTTPS or OAuth2 certificate on one or more devices..

Monitoring XCO Status

The Monitoring service provides REST API to monitor the status of the various services running in XCO.

The service runs on the host and is exposed on port 8078, which is not the port where the XCO application is running. In a multi-node deployment, this service is available on both nodes and can be accessed through the virtual IP (VIP).

- To start or stop the Monitoring service, run the **`systemctl stop/start/restart monitor.service`** command as a sudo or root user.

  For information, see REST API Documentation for XCO on page 23.

- Use the **`efa status`** command to verify application status.

  For more information, see the **`efa status`** command in the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Verifying XCO System Health

Use this topic to learn about the methods for verifying the health of various XCO services.

SLX Device Health

By default, health check functionality is deactivated when SLX devices are registered. You can verify the status of the functionality with the following XCO command.

```
$ efa inventory device setting show --ip <ip-addr>

+--------------------------------+-------+
| NAME                           | VALUE |
+--------------------------------+-------+
| Maintenance Mode Enable On     | No    |
| Reboot                         |       |
+--------------------------------+-------+
| Maintenance Mode Enable        | No    |
+--------------------------------+-------+
| Health Check Enabled           | No    |
+--------------------------------+-------+
| Health Check Interval          | 6m    |
+--------------------------------+-------+
| Health Check Heartbeat Miss    | 2     |
| Threshold                      |       |
+--------------------------------+-------+
| Periodic Backup Enabled        | Yes   |
+--------------------------------+-------+
| Config Backup Interval         | 24h   |
+--------------------------------+-------+
| Config Backup Count            | 4     |
+--------------------------------+-------+
--- Time Elapsed: 270.251797ms ---
```

You can enable health check functionality on the device. And you can configure XCO to regularly back up the device configuration (every 6 minutes by default). For more information, see Configure Backup and Replay on page 460.

If the threshold for missed heartbeats is exceeded, XCO begins the drift and reconcile process after connectivity to the device is re-established. For more information, see Drift and Reconcile on page 60.

XCO Services Health

All services in XCO have internal health REST APIs that Kubernetes uses to restart pods that are deemed unhealthy. The results of a liveness probe determines whether a pod is healthy. Typical values for liveness probes are as follows:

- initialDelaySeconds: 60
- periodSeconds: 10

- timeoutSeconds: 15

RabbitMQ Liveness

The XCO message bus is the workhorse for asynchronous inter-service communication. Therefore, XCO uses the RabbitMQ built-in ping functionality to determine the liveness of the RabbitMQ pod.

As part of a health check, each XCO service also validates its connection to RabbitMQ and attempts to reconnect to RabbitMQ when necessary.

XCO System Health for High-availability Deployments

During installation or upgrade of XCO, a system service called `efamonitor` is set up. This service runs validations every minute to check XCO database cluster, glusterFS, and RabbitMQ are functioning correctly.

As needed, the `efamonitor` service remediates the MariaDB Galera cluster and RabbitMQ connection issues, and logs the stats of the system.

Node Health

To ensure that the active and standby nodes are operational, ping checks occur between the nodes. The pings determine whether the active node is up and running. If not, the virtual IP addresses are switched over to the other node.

To ensure that failover does not occur due to a network issue, if a ping to the peer fails, a ping is also attempted to the default gateway. If ping to default gateway fails, ping is attempted to any alternative gateway that may have been provided during installation or upgrade.

If all of the pings fail, keepalived triggers Kubernetes to switch over to the active node and to put the other node in a Fault state.

XCO System Backup and Restoration

The backup process saves XCO data, including the database, certificates, and multi-access network configuration. The process does not back up northbound certificates.

Manual Backup and Restore

XCO supports backup and restore across the same IP modes. A warning message appears during restore operation.

The following table describes the support matrix for backup and restore operation across the IP modes:

Backup	Restore	Support
Dual	IPv6	No
Dual	IPv4	No
IPv4	IPv6	No
IPv4	Dual	Yes
IPv6	IPv6	No
IPv6	IPv4	No
IPv4	IPv4	Yes
IPv6	IPv6	Yes
Dual	Dual	Yes

The backup process creates a backup tar file. You can select from all saved tar files during the restore process. The tar files are saved to one of the following locations.

- Server: `/var/log/efa/backup`
- TPVM: `/apps/efa_logs/backup`

A backup generated on one XCO system can be restored on another system.

For more information, see Back up and Restore the XCO System on page 50.

You can use the **`efa system backup-list`** command to see the backup files that are available to use in a restore operation. For example:

```
$ efa system backup-list
+----+-----------------------------------+---------+----------------+
| ID    | File                           | Version | Generated By   |
+----+-----------------------------------+---------+----------------+
| 1     | EFA-2020.08.20-20.26.46.tar    | 2.3.0-1 | User           |
+----+-----------------------------------+---------+----------------+
| 2     | EFA-2020.08.20-20.27.29.tar    | 2.3.0-GA| System         |
+----+-----------------------------------+---------+----------------+
--- Time Elapsed: 183.69386ms ---
```

Periodic Backups and Configuration

XCO periodically backs up the system (12 AM by default, but configurable). During the backup process, all the services are locked and APIs return the message "Service is Locked with reason backup". Therefore, you cannot see the location under dropdown for Device Discovery.

The periodic backup process creates a backup tar file that is saved to the same location as the manual backup files. When new backup and supportsave files are created, according to the age of the backup files and maximum number of backup files to save, the system deletes the saved system-generated backup files, supportsave files, and manual backup files.

You can use the **`efa system settings update`** command to determine the backup schedule and to change the maximum number of backup files to save.

- Default backup schedule: `0:*:*:*`, meaning every day at midnight.
- Default maximum number of backup files: Five backup files and five supportsave files.

For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

You can use the **`efa system settings show`** command to view the current backup settings.

Tip
You can use the **`efa system cleanup`** REST API to delete a specified backup or supportsave file. This feature lets you delete files before they are automatically deleted.

You can use the **`efa system settings reset`** command to reset the backup system settings to default values.

Backups During Upgrades

The upgrade process of XCO also backs up the XCO system to recover data if the upgrade fails. For more information, see "Recover from an Upgrade Failure" in the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

Logs

Logs for backup, restore, and supportsave operations are saved to the following locations:

- Server: `/var/log/efa/system` and `/var/log/efa/monitor`
- TPVM: `/apps/efa_logs/system` and `/apps/efa_logs/monitor`

The REST APIs for backup, restore, and supportsave are part of the Monitoring Service and can be accessed via port 8078. The logs for the REST APIs are saved to the `<log_dir>/monitor/` location.

Back up and Restore the XCO System

You can back up and restore the XCO system, including the database and certificates.

Procedure

1. To back up the system, run the following command.

```
$ efa system backup
Generating backup of EFA...
Backup Location: /apps/efa_logs/backup/EFA-2021-02-10T13-21-46.413.tar
--- Time Elapsed: 10.401999384s ---
```

2. To restore the system, take the following steps.
   a. Run the **efa system restore** command.

```
$ efa system restore

EFA-2021.01.12-11.19.52.tar (Version:2.3.2-GA, Generated by: User)
EFA-2021-01-12T04.59.09.tar (Version:2.4.0-7171, Generated by: User)
EFA-2021-01-12T13.50.00.tar (Version:2.4.0-121, Generated by: User)
EFA-2021-01-12T13.50.51.tar (Version:2.4.0-121, Generated by: System)
EFA-2021-01-12T16.11.47.tar (Version:2.4.0-1211, Generated by: System)
EFA-Upgrade-2.3.2-GA.tar (Version:2.3.2-GA, Generated by: System)
```

   The command output displays a list of available backup tar files.
   b. Select the backup tar file that you want to restore.

```
Choose backup option:1
Selected: EFA-2021.01.12-11.19.52.tar
Performing EFA restore using EFA-2021.01.12-11.19.52.tar
Generating backup before initiating restore
BACKUP_TAR: /apps/efa_logs/backup/EFA-2021.01.12-11.19.52.tar
Stopping all EFA services
All pods are terminated
Migrating database
Completed database migration
Checking if all PODS are in ready state...
Restore operation is successful
--- Time Elapsed: 9m3.079104969s ---
```

   c. When the restore is complete, run **source /etc/profile**.

      You can now log in to XCO.
   d. To enable secure connections, install the certificates on devices.

```
efa certificate device install --ip 10.20.61.92 --cert-type https
```

      The command installs the HTTPS or OAuth2 certificate on one or more devices.
   e. To get the current state of the devices, run the **efa inventory device update** command after you run the restore command.
   f. Check the status of the services to ensure that they are in-sync.

Change the Host Name or IP Address

You can change the host name, the IP address, and the virtual IP address after XCO is deployed.

Before You Begin

Review the following host name requirements:

- Host name changes are supported in single-node and multi-node deployments.
- IP address changes are supported in single-node deployments.
- Virtual IP address (VIP) changes are supported in multi-node deployments.
- Host names must be unique and consist of numeric characters and lowercase alphabetic characters. Do not use uppercase alphabetic characters.
- Hyphens are the only special characters allowed. No other special characters are allowed by Kubernetes for cluster formation or by the K3s service.

About This Task

Follow this procedure to change the host name, IP address, and virtual IP address.

Procedure

1. To change the host name, take the following steps.
   a. On a server installation, run the following Linux command to change the host name of the system:

      ```
      hostnamectl set-hostname <new name>
      ```

      - Update the new host name in `/etc/hosts`.
      - In a TPVM deployment, run the following SLX command to change the host name of the system.

        ```
        device(config-tpvm-TPVM)# hostname <new name>
        ```

   b. Run the following command as a root user or as a user with sudo privileges.

      Important
      Do not reboot the system before running this command.

      ```
      sudo bash efa-change-hostname <old host name>

      Reading host name of the system
      Restarting mariadb service
      Restarting k3s service
      Checking k3s for the new host name
      Host is in ready state in k3s
      Setting current host as active node
      Deleting old host name references
      Waiting for EFA containers to start
      Successfully updated host name in EFA
      ```

      In a single-node deployment, XCO is not operational during this step. In a multi-node deployment, XCO remains operational if the command is running on the standby node. XCO is not operational if the command is running on the active node.
2. To change the IP address of a single-node deployment, take the following steps.
   a. Run the following command as a root user or as a user with sudo privileges.

      ```
      sudo bash efa-change-ip

      Updating IP in EFA
      Restarting k3s service
      Updating all files with new IP
      Deleting EFA services: gonotification-service gofabric-service gotenant-service
      goauth-service gorbac-service goinventory-service goopenstack-service
      govcenter-service gohyperv-service goraslog-service efa-api-docs gosystem-service
      Waiting for EFA containers to start
      Successfully updated IP in EFA
      ```

      XCO is not operational during this step.

      In a TPVM deployment, you can run the command from `/apps/bin/`.

      In a single-node deployment of XCO on TPVM, changing the IP address of a node is not supported for all the IP mode (IPv4, IPv6, and Dual IP modes) deployment types.
   b. After the IP address is updated, run `source /etc/profile` or open a new XCO session to log in.

3. To change the VIP of a multi-node deployment, take the following steps.
   a. Stop the monitoring service.

```
sudo systemctl stop efamonitor
```

   b. Start the monitoring service.

```
sudo systemctl start efamonitor
```

Note
After the Virtual IP address (VIP) update is successful, it takes a few minutes to update new VIP in all the registered devices in XCO.

Display XCO Running Configurations

You can view the running-config of all current XCO configurations for core services.

About This Task

The output is displayed in the following order: Asset, Fabric, Tenant commands, Policy, and System running configuration. The command output contains the default values for each configuration line item.

You can use the command output for CLI playback on an empty XCO deployment, which is a useful tool for recovery.

Note
The output of **`efa show-running-config`** command is also captured as part of the supportsave zip file.

Procedure

Run the **`efa show-running-config`** command.

```
$ efa show-running-config

efa inventory device register --ip "10.24.80.191" --username admin --password password

efa inventory device setting update --ip "10.24.80.191" --maint-mode-enable-on-reboot  No
--maint-mode-enable No --health-check-enable No --health-check-interval 6m
--health-check-heartbeat-miss-threshold 2 --config-backup-periodic-enable Yes
--config-backup-interval 24h --number-of-config-backups 4

efa inventory device register --ip "10.24.80.192" --username admin --password password

efa inventory device setting update --ip "10.24.80.192" --maint-mode-enable-on-reboot  No
--maint-mode-enable No --health-check-enable No --health-check-interval 6m
--health-check-heartbeat-miss-threshold 2 --config-backup-periodic-enable Yes
--config-backup-interval 24h --number-of-config-backups 4

efa fabric create --name "default" --type clos --stage 3 --description "Default Fabric"
```

This example shows only a partial list of typical output.

Audit Trail Logging

XCO provides full audit trail logging, including the successes and failures of user actions, which creates a 1-to-1 mapping between every action coming from XCO and a corresponding audit trail event from SLX.

Any configuration action on an SLX devices results in the generation of an audit trail. The name of the user is extracted from the token that the user logged in with. The user is assigned the role of `admin` as the default role on the device.

For OpenStack, the user name has the following format: `<OpenStack tenant UUID> - <OpenStack user name> - <EFA tenant name>`.

The following is an example of the audit log message for NETCONF or SSH sessions:

```
78 AUDIT, 2020/01/26-14:04:21 (GMT), [DCM-1006], INFO, DCMCFG, <ClientUserID>/
<ClientRole>/10.6.46.51/SSH/netconf,, SLX, Event: database commit transaction, Status:
Succeeded, User command: "configure config username test1 role admin password ****".
```

The `ClientUserID` and `ClientRole` values are derived from the `User` and `AuditLogRole` variables, which originate from the values in the access token when the NETCONF or SSH session was established.

Transfer of Audit Trail Data

Audit trail data from SLX devices is transferred to XCO for delivery upstream using JSON structured data.

The data is transferred to an upstream web server at a predefined URL that is registered with XCO.

Incoming syslog messages from SLX to XCO are converted by a logging service on XCO into JSON data, as in the following example:

```
"message_id": "9999",
  "message": "Hello world",
  "source_ip": "192.168.10.1",
  "user": "admin",
  "severity": "INFO",
  "timestamp", "2020-02-11 19:23:58.383304",
  "extra_data": {}
}
```

XCO sends the messages by POST requests to an upstream web receiver.

Logging and Log Files

XCO logs are saved to the following locations:

- Non-TPVM deployments: `/var/log/efa`. The installation logs in the `/var/log/efa/installer` directory are a good source for discovering the reason for a failure.
- TPVM deployments: `/apps/efa_logs`

- Kubernetes log files: `/var/log/pods`
- Keepalived service log files: `/var/log/keepalived`. The directory contains Keepalived service logs and the `journalctl` log of Keepalived for the past day. This information is useful for helping debug failovers, double faults, and gateway connectivity.

In multi-node, high availability deployments, logs are replicated on all nodes in the cluster.

The **`efa system supportsave`** script gathers all logs, database dumps, pod logs, deployment details, and system support-save and then compresses them into a ZIP folder. You can share this ZIP folder with Extreme support personnel when troubleshooting an issue.

Logging Customization

You can limit the log types that are stored and customize the log files to fit into the system resources. Systems with high resources can hold more logs for a longer period of time. Systems with low resources can reduce logging to fit the system needs.

Each service has two different customizable logs: all log, and the error log. Each customizable log may have up to 10 files of history: 1 active log file and 9 compressed history files. The active log file has a maximum size of 100MB for the base log, and 10MB for the error log. The default logging configuration for the base log captures all logging levels between panic and debugging, while the error log captures all logging levels between panic and error. You can customize all these values through logging customization.

Note
Use the **`efa system logging show`** command to view the current logging configuration.

Configure Logging

You can customized logging of a service. By default, logging is configured with its default parameters and no logging types are excluded. Any log entry exceeding set limits will be lost.

About This Task

Follow this procedure to configure a customized logging.

Procedure

To configure the logging of a service to a customized state, run the **`efa system logging set`** command with appropriate parameters.

```
efa system logging set --service inventory --size 200 --level info
+-----------+------+------+---------------+-------+---------+--------+
|  Service  | Type | Size | Maximum Files | Level | Status  | Reason |
+-----------+------+------+---------------+-------+---------+--------+
| inventory | all  | 200  |               | info  | Success |        |
```

```
+-----------+------+------+---------------+-------+---------+--------+
Logging set

efa system logging set --service inventory,fabric --type error --maximum-files 20
+-----------+-------+------+---------------+-------+---------+--------+
|  Service  | Type  | Size | Maximum Files | Level | Status  | Reason |
+-----------+-------+------+---------------+-------+---------+--------+
| inventory | error |      |      20       |       | Success |        |
+-----------+-------+------+---------------+-------+---------+--------+
|  fabric   | error |      |      20       |       | Success |        |
+-----------+-------+------+---------------+-------+---------+--------+
Logging set

efa system logging set --service inventory --type error --level fatal
+-----------+-------+------+---------------+-------+---------+--------+
|  Service  | Type  | Size | Maximum Files | Level | Status  | Reason |
+-----------+-------+------+---------------+-------+---------+--------+
| inventory | error |      |               | fatal | Success |        |
+-----------+-------+------+---------------+-------+---------+--------+
Logging set
```

Note

For information about commands and supported parameters to configure logging, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Unconfigure Logging

You can unconfigure a logging customization. The unconfiguration sets all the logging configurations to their default values.

About This Task

Follow this procedure to unconfigure a customized logging.

Procedure

To unconfigure logging customization of a service to the default state, run the following command:

```
efa system logging unset [ --service service-name | --type logging-type |
```

Note

For information about commands and supported parameters to unconfigure logging customization, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Example

The following example sets the logging configuration to default:

```
efa system logging unset --service inventory
+-----------+------+---------+--------+
|  Service  | Type | Status  | Reason |
+-----------+------+---------+--------+
| inventory | all  | Success |        |
+-----------+------+---------+--------+
Logging unset

efa system logging unset --service inventory,fabric --type error
```

```
+-----------+-------+---------+--------+
|  Service  | Type  | Status  | Reason |
+-----------+-------+---------+--------+
| inventory | error | Success |        |
+-----------+-------+---------+--------+
|  fabric   | error | Success |        |
+-----------+-------+---------+--------+
Logging unset
```

Data Consistency

XCO ensures that SLX devices have the correct configuration before allowing traffic.

Overview

XCO is the data owner and Single Source of Truth (SSOT) for fabric configuration. The following figure describes how data is rendered consistent among XCO services.

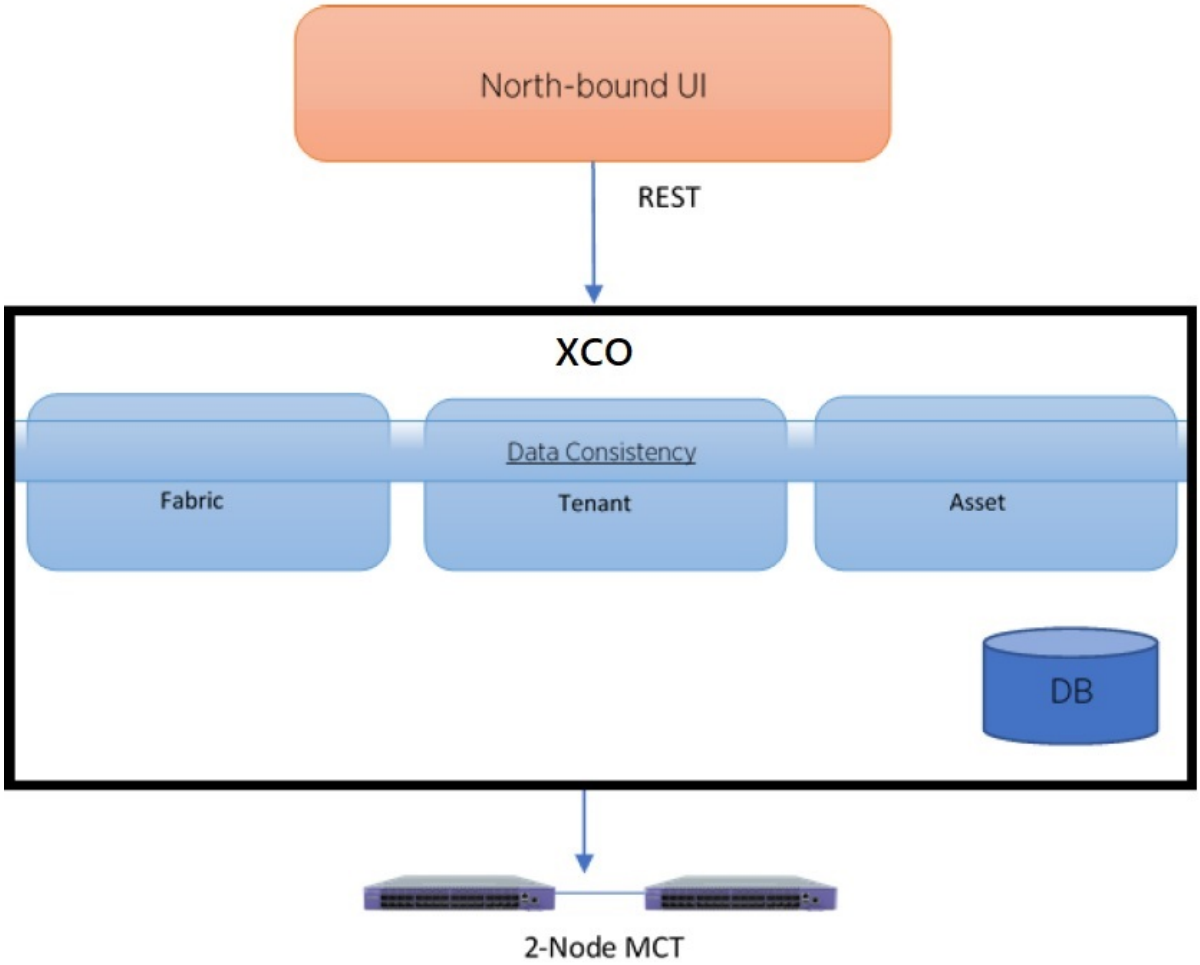

Figure 3: Data consistency overview

North-bound applications invoke REST APIs to perform various operations on XCO. XCO ensures that the operations leave XCO and the fabric in a consistent state.

Limitations

- You cannot use the SLX CLI to configure the entities that are managed by XCO.
- XCO can reconcile only those entities or configurations that it manages.
- XCO cannot modify out-of-band entities or configurations unless they conflict with the configurations that it manages.

Periodic Device Discovery

Asset, Tenant, Fabric, and Policy Services use periodic discovery to detect out-of-sync configurations on the devices. These Services act on the published events and update the database to reflect the status of the devices as in-sync and out-of-sync.

You can perform on-demand full device discovery using the **`efa inventory device update`** command.

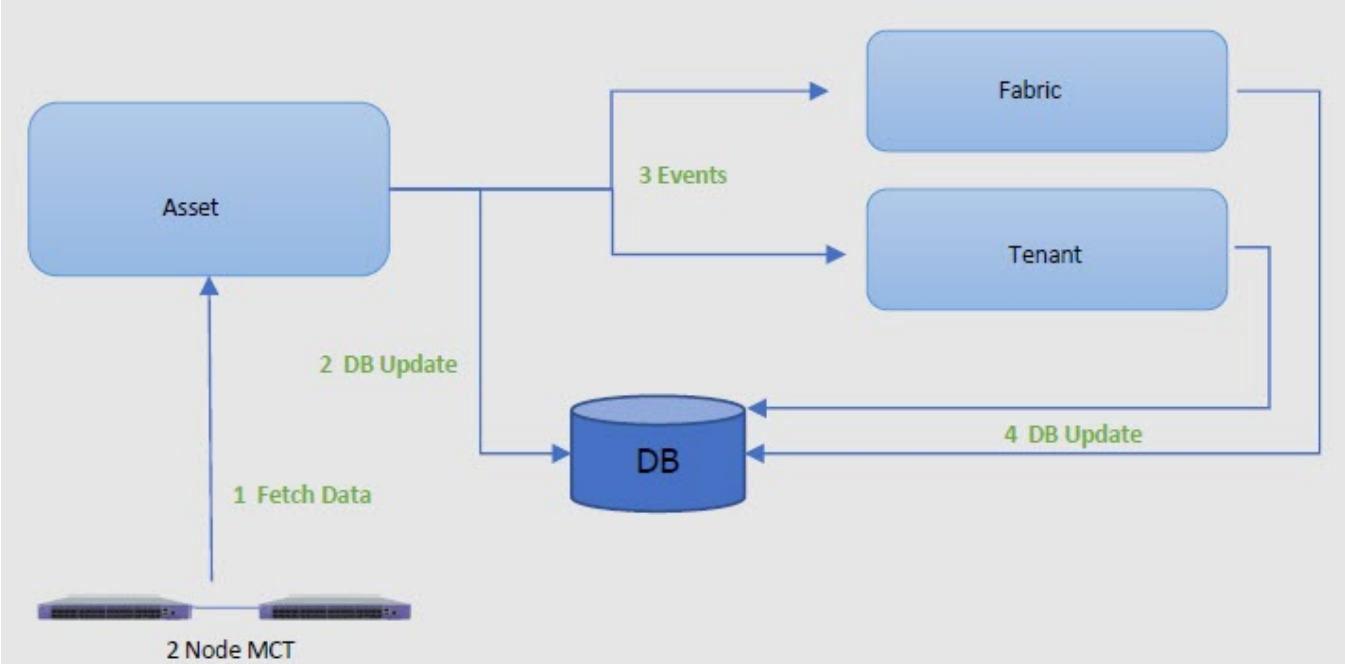

Figure 4: Device discovery and database updates

The Asset service periodically polls the devices in the fabric and keeps the database and other services updated of any changes in the underlying fabric. The default polling interval is one hour, with valid values ranging from 15 minutes to 24 hours.

You can use the **`efa inventory device discovery-time list`** command to view the current discovery interval for a device or fabric. You can use the **`efa inventory device discovery-time update`** command to configure the discovery interval.

XCO determines out-of-band configuration changes on the devices. If there are no out-of-band configuration changes, the device updates are optimized.

Note

- Periodic device discovery is compatible with SLXOS 20.4.1 and later.
- Log entries do not get populated for the device discovery failures. Only the success cases get listed under the Device Discovery.

Persistent Configuration

Extreme devices support three types of configuration files:

- Default - Default configuration files are part of the firmware package for the device and are automatically applied to the startup configuration.
- Startup - Startup configuration files are persistent and are applied after system reboot.
- Running - Configuration currently effective on the device is the running configuration.

For more information on configuration files, see Extreme SLX-OS Management Configuration Guide.

In SLX-OS 20.1.1, the configuration management process maintains two databases, Running and Startup.

In SLX-OS 20.1.2 and later, all the configurations are stored in one database, which also persists.

- The **`show running-config`** command fetches the configuration from the database.
- The **`copy running-config startup-config`** command creates or updates the persistent configuration.
- After a upgrade or downgrade, replaying the startup file resumes the SLX database cleanup operations.

Maintenance Mode

In SLX-OS 20.1.1, maintenance mode can be enabled by configuring **`enable`** under system-maintenance configuration mode. If the configuration is persistent, the switch needs to be in maintenance mode before rebooting for it to come back in maintenance mode.

In SLX-OS 20.2.1 and later, maintenance mode can be enabled by configuring **`enable-on-reboot`** under system-maintenance configuration mode. After the reboot, the device comes back up in maintenance mode and remains operational.

```
SLX(config-system-maintenance)# enable-on-reboot
SLX(config-system-maintenance)# [no] enable-on-reboot
```

The **`system maintenance turn-off`** command brings the system out of maintenance mode.

Non-reachable Devices

XCO tracks devices by running heart-beats to the SLX devices.

When a non-reachable device becomes reachable, XCO identifies any drift and performs reconciliation, if necessary.

Drift and Reconcile

XCO supports drift and reconcile (DRC) of a configuration at device level. A single device configuration is compared with XCO. If there is a drift in the configuration, it is conciled. XCO provides APIs to initiate drift and reconcile requests. Use the XCO command **`efa inventory drift-reconcile execute`** to run a manual DRC.

Drift and reconcile is also activated during the following operations:

- Switch replacement and RMA
- After the reboot of a device in maintenance mode
- Device firmware-download with “drc” flag

Drift and reconcile operations are run in parallel across all devices in the fabric. It ensures that the multiple DRC operations that take place during fabric-wide firmware-download (FWDL) or reboot of multiple devices together, run in parallel, and hence, reduce the overall maintenance window.

Note

If **`maintenance-mode-enable on reboot`** is not set on the devices, Data Consistency is not guaranteed and drift and reconciliation operation is skipped.

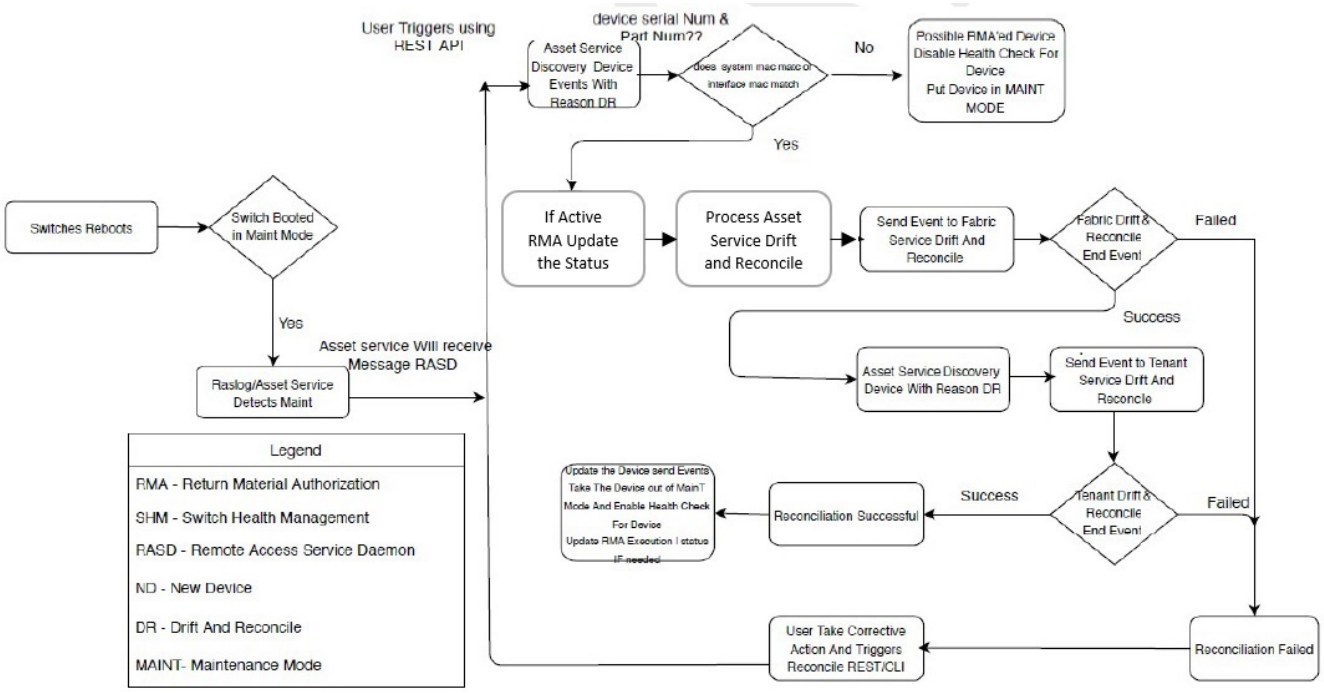

Figure 5: Drift and reconcile workflow

Note

When any attribute under "router bgp" is drifted, XCO also reconciles the cluster configuration to ensure that the BGP neighbors of MCT are reconciled, and this shows up as cluster reconciled success in addition to routerbgp.

Network Elements

Starting in XCO v2.5.0, in addition to fabric and tenant service configurations, the following asset service configurations are persisted and included in Drift and Reconcile (DRC).

The support is on two levels:

- **Interface level configuration**: Breakout mode, MTU, admin state, speed, FEC configuration, port dampening (link-error disable), and RME
- **Global or system level configuration**: NTP, SNMP v2 and v3, prefix list, and route map

Interface-level Configuration

The following table captures the various attributes of interface for which DRC and idempotency is supported.

- A drift is identified if any of the fields below is modified through the SLX, CLI command, or other management tool.

- A reconcile operation pushes the intended configuration to SLX, so bringing the SLX configuration in sync with XCO.

Note
Clean up explicitly any conflicting configuration which could cause reconciliation of device to fail. For example, if XCO configures a port as breakout and if that configuration is drifted by adding Layer 3 configuration to a parent interface, the reconciliation fails. It is recommended to explicitly remove the conflicting configuration from the device through the SLX CLI and retry the DRC process.

Table 5: Interface attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile configuration	Idempotency
Admin-state	Yes	Yes	Yes
Breakout mode	Yes	Yes	No
Speed	Yes	Yes	Yes
Layer-2 MTU	Yes	Yes	Yes
IPv4 MTU	Yes	Yes	Yes
IPv6 MTU	Yes	Yes	Yes
FEC mode	Yes	Yes	Yes
Link error disable	Yes	Yes	Yes
Toggle-threshold	Yes	Yes	Yes
Sampling time	Yes	Yes	Yes
Wait time	Yes	Yes	Yes
RME enable	Yes	Yes	Yes

The following CLI commands are available:

- `efa inventory device interface redundant-management`
- `efa inventory device interface set-fec`
- `efa inventory device interface set-link-error-disable`
- `efa inventory device interface unset-fec`
- `efa inventory device interface unset-link-error-disable`

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Global or System-level Configuration

- A drift is identified if any of the fields below is modified through the SLX, CLI command, or other management tools.
- A reconcile operation pushes the intended configuration to SLX, so bringing the SLX configuration in sync with XCO.

The following CLI commands are available:

- `efa inventory device setting update --prefix-independent-convergence`
- `efa inventory device setting update --prefix-independent-convergence-static`
- `efa inventory device setting update --maximum-load-sharing-paths`
- `efa inventory device settings update --mct-bring-up-delay`
- `efa inventory device settings update --maint-mode-convergence-time`

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Device Reload

The device reload command allows the user to reload a device. Users can provide IPs and fabric name separated by commas. All devices in a fabric will be reloaded for any given fabric name in the command.

The following CLI command is available for device reload:

- `efa inventory device reload`

Note
Drift and reconcile and idempotency configuration support is not applicable for device reload attribute.

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Clear IP Route

The clear IP route command allows the user to clear a device's IPv4 and IPv6 routes. Users have the option to either clear an IPv4 or an IPv6 route.

The following CLI command is available for clearing IP route:

- `efa inventory device clear route-all`

Note
Drift and reconcile and idempotency configuration support is not applicable for clear IP route attribute.

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

SNMP Configuration

The following tables capture the various attributes of SNMP and NTP for which DRC and idempotency are supported.

Note
Regarding idempotency for creating an entry which already exists in XCO, an error message is returned stating that the user already exists.

Table 6: SNMP attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile Configuration	Idempotency	Notes
Community deleted	Yes	Yes	No	A valid error message is shown when a non-existent community is deleted
Group name associated with community is modified	Yes	Yes	Not Applicable	
Group deleted	Yes	Yes	Not Applicable	
Modify group version	No	No	Not Applicable	SLX does not support editing the SNMP group version
Modify read review or write view or notify view associated with group	Yes	Yes	Not Applicable	
Modify group name associated with SNMP user	Yes	Yes	Not Applicable	
Modify authentication protocol associated with SNMP user	Yes	Yes	Not Applicable	
Modify authentication password associated with SNMP user	Yes	Yes	Not Applicable	
Modify privacy protocol associated with SNMP user	Yes	Yes	Not Applicable	

Table 6: SNMP attributes supporting DRC and Idempotency (continued)

Field	Identify Drift	Reconcile Configuration	Idempotency	Notes
Modify privacy password associated with SNMP user	Yes	Yes	Not Applicable	
Delete SNMP user	Yes	Yes	No	A valid error message is shown when a non-existent user is deleted
Modify encrypted keyword associated with SNMP user	Yes	Yes	Not Applicable	
Modify authentication type associated with group, that is, auth, noauth, notify	Yes	Yes	Not Applicable	
Delete SNMP host entry	Yes	Yes	No	A valid error message is shown when a non-existent host is deleted
Update SNMP host severity level	Yes	Yes	Not Applicable	Any drift observed from XCO configured default severity level is reconciled
Update SNMP host source interface	Yes	Yes	Not Applicable	Any drift observed from XCO configured default source interface is reconciled
Update SNMP host UDP port	Yes	Yes	Not Applicable	Any drift observed from XCO configured default UDP port is reconciled
Update SNMP host VRF	Yes	Yes	Not Applicable	Any drift observed from XCO configured default VRF is reconciled
Update SNMP host engine id	Yes	Yes	Not Applicable	

Table 6: SNMP attributes supporting DRC and Idempotency (continued)

Field	Identify Drift	Reconcile Configuration	Idempotency	Notes
Update of SNMP host notification type (traps, informs)	Yes	Yes	Not Applicable	
Update of SNMP view MIB OID access (included, excluded)	Yes	Yes	Not Applicable	

The following CLI commands are available for operations on SNMP interfaces:

- `efa inventory device snmp community create`
- `efa inventory device snmp community delete`
- `efa inventory device snmp community list`
- `efa inventory device snmp user create`
- `efa inventory device snmp user delete`
- `efa inventory device snmp user list`
- `efa inventory device snmp host create`
- `efa inventory device snmp host delete`
- `efa inventory device snmp host list`

For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

NTP Configuration

The NTP commands let you configure NTP server configuration on the SLX device. The configuration you set is persisted in the XCO database. DRC is also supported.

The following table captures the various attributes of the NTP configuration interface for which DRC and idempotency is supported. A drift is identified if any of the following fields are modified by you through SLX CLI or other management tools. Reconcile operation pushes the intended configuration to SLX which makes the SLX configuration synchronize with XCO.

On idempotency for creating an entry which already exists in XCO, an error message is returned stating that user already exists.

Table 7: NTP attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile configuration	Idempotency	Notes
NTP auth key ID associated with NTP serer is modified	Yes	Yes	Not Applicable	
NTP auth key name associated with NTP serer is modified	Yes	Yes	Not Applicable	
NTP server deleted	Yes	Yes	No	A valid error message is shown when a non existent NTP server is deleted.
Encryption type is modified	Yes	Yes	Not Applicable	
Trusted key is modified	Yes	Yes	Not Applicable	
Encryption level is modified	Yes	Yes	Not Applicable	
NTP server disable modified	Yes	Yes	Not Applicable	

The following CLI commands are available for operations on NTP interfaces:

- `efa inventory device ntp server create`
- `efa inventory device ntp server delete`
- `efa inventory device ntp server list`
- `efa inventory device ntp disable-server`

For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

IP prefix list configuration

Note
Regarding idempotency for creating an entry which already exists in XCO, an error message is returned stating that the user already exists.

Table 8: IP prefix list attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile configuration	Idempotency	Notes
IPv4 prefix list rule is deleted.	Yes	Yes	No	Deleted rule will be reconciled.
IPv4 prefix list is deleted.	Yes	Yes	No	Deleted prefix list along with all rules associated with it will be reconciled.
IPv4 prefix list rule created OOB. Different rules exist with same prefix list name in XCO.	No	No	Not applicable	Delete the OOB rule or keep it and don't act as part of DRC.
IPv4 prefix list rule created OOB. Different rules exist with same prefix list name and sequence number in XCO.	Yes	Yes	Not applicable	Prefix list rule will be reconciled to be in sync with XCO.
Create an IPv4 prefix OOB with a prefix list name not matching any of the XCO created entries.	No	No	Not applicable	These are treated as out of band entries and XCO will not perform DRC.

The following CLI commands are available for operations on IP prefix lists:

- **`efa policy prefix-list create`**
- **`efa policy prefix-list list`**
- **`efa policy prefix-list delete`**
- **`efa policy prefix-list update`**

For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Route map configuration

Note
Regarding idempotency for creating an entry which already exists in XCO, an error message is returned stating that the user already exists.

Table 9: Route map attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile configuration	Idempotency	Notes
Route map deleted.	Yes	Yes	No	Recreate the route map, along with the match criteria during reconcile.
Route map rule action updated.	Yes	Yes	No	Reconcile the route map action (permit/deny) for that rule.
Update IPv4 prefix list name in match criteria.	Yes	Yes	No	Reconcile the IPv4 prefix list name.
IPv4 prefix list match criteria deleted.	Yes	Yes	NA	Reconcile the match criteria for IPv4 prefix list.
A different match criteria NOT supported by XCO is added through OOB.	No	No	NA	
A set criteria NOT supported by XCO is added through OOB.	No	No	NA	
Route map is created through OOB and this is not present or created by XCO.	No	No	NA	

The following CLI commands are available for operations on route maps:

- **`efa policy route-map create`**
- **`efa policy route-map update`**
- **`efa policy route-map delete`**
- **`efa policy route-map list`**
- **`efa policy route-map-match create`**
- **`efa policy route-map-match list`**
- **`efa policy route-map-match delete`**

For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Device Settings

The following table captures the various attributes of device settings for which DRC and idempotency are supported.

Table 10: Device settings attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile Configuration	Idempotency
BGP prefix independent convergence (PIC)	Yes	Yes	Yes
prefix-independent-convergence-static	Yes	Yes	Yes
ECMP routed load-sharing max path	Yes	Yes	Yes
Maintenance mode convergence time	Yes	Yes	Yes
Static prefix independent convergence (PIC)	Yes	Yes	Yes

Note
Drift and reconcile and idempotency configuration support is not applicable for device update and viewing device settings.

BGP Prefix Independent Convergence (PIC)

Specify **`Yes`** to enable BGP PIC and **`No`** to de-configure it.

The following CLI command is available to enable BGP PIC:

- `efa inventory device setting update --prefix-independent-convergence`

Note
After configuring this command, clear the routes on the device.

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

ECMP Max Path

Use the command **`string`** to view route load-sharing maximum paths. Valid values include **`8`**, **`16`**, **`32`**, and **`64`**. The default value is **`64`**.

Note
The device must be reloaded for this command to take effect.

The following CLI command is available to configure ECMP route load-sharing max path:

- `efa inventory device setting update --maximum-load-sharing-paths`

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Device Settings Update

Configure the maintenance mode and display the available device settings.

Use the **`device-ips`** parameter, separated by comma to view a range of device IP addresses.

Use the **`fabric-name`** parameter to specify the name of the fabric.

The **`show`** command displays the device settings.

The following CLI command is used to display the device settings:

- `efa inventory device setting show [ --ip device-ips ]`

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Time Zone Configuration

By default, SLX devices come up with the GMT timezone. Using the `efa inventory device timezone` command, you can set the timezone per device or per fabric.

The following CLI commands are available for timezone settings:

- **`efa inventory device timezone set`**
- **`efa inventory device timezone unset`**
- **`efa inventory device timezone list`**

Table 11: Time zone attributes supporting DRC and Idempotency

Field	Identify Drift	Reconcile configuration	Idempotency
Time zone is set.	Yes	Yes	Yes
Time zone is unset.	Not applicable	Not applicable	Yes

Note
Identify drift, drift and reconcile, and idempotency support is not applicable for time zone display.

For more information, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Idempotent Operations

The idempotent operations produce the same result for multiple identical requests or operations.

Reissuing an XCO command should leave the system in the same state as the last time the command was run. Such idempotent operations help ensure data consistency during high-availability failovers.

In this example, running the **`efa fabric create`** command twice, with the same parameters, produces the same result each time.

```
$ efa fabric create --name fabric1  --type non-clos --description non-clos-fabric
Create Fabric nonclos [Success]

(efa:extreme)extreme@tpvm:~$ efa fabric create --name fabric1 --type non-clos --
description non-clos-fabric
Create Fabric nonclos [Success]
```

Rollback Scenarios for Data Consistency

Rollback of failed configuration changes ensures data consistency.

Failure on Some Devices during Configuration

When a REST operation succeeds on one device but fails on another, configuration changes are rolled back for both devices. In the following example, the operation fails on one MCT node but succeeds on the other. The whole operation fails and an error message is returned as part of the REST response.

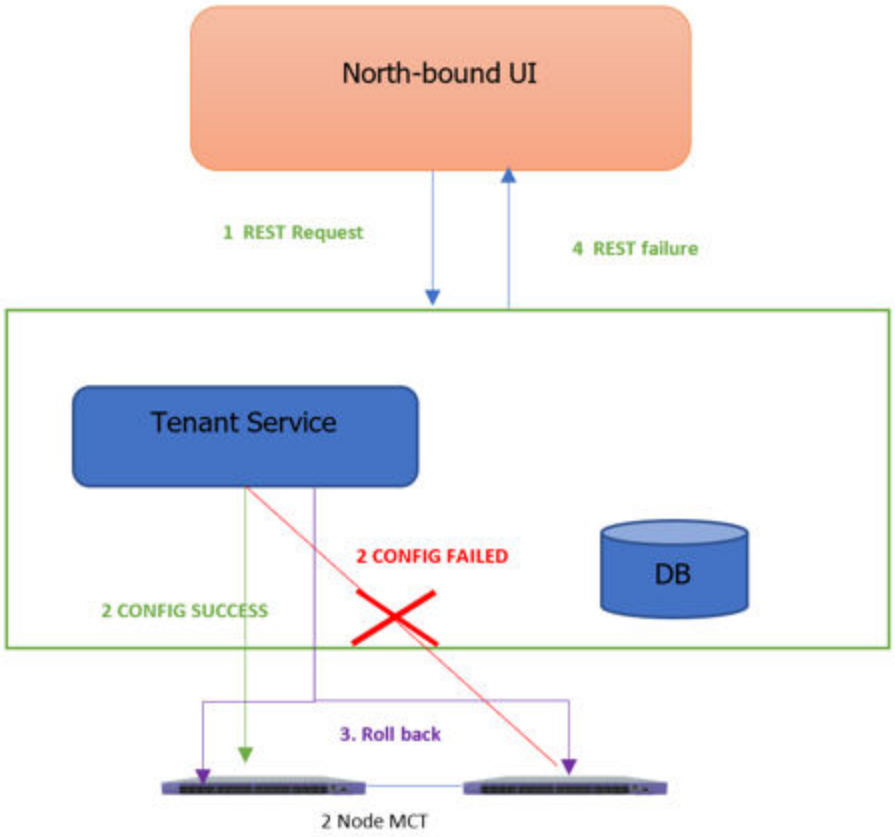

Figure 6: Rollback for failure of one node

Note
This process for partial failures is the default. You can change the process to enable partial successes even when one node fails. For more information, see Administered Partial Success on page 362.

Failure on All Devices during Configuration

When a REST operation fails on all devices in the request, configuration changes are rolled back for all devices. In this example, the operation fails on both MCT nodes and an error message is returned as part of the REST response.

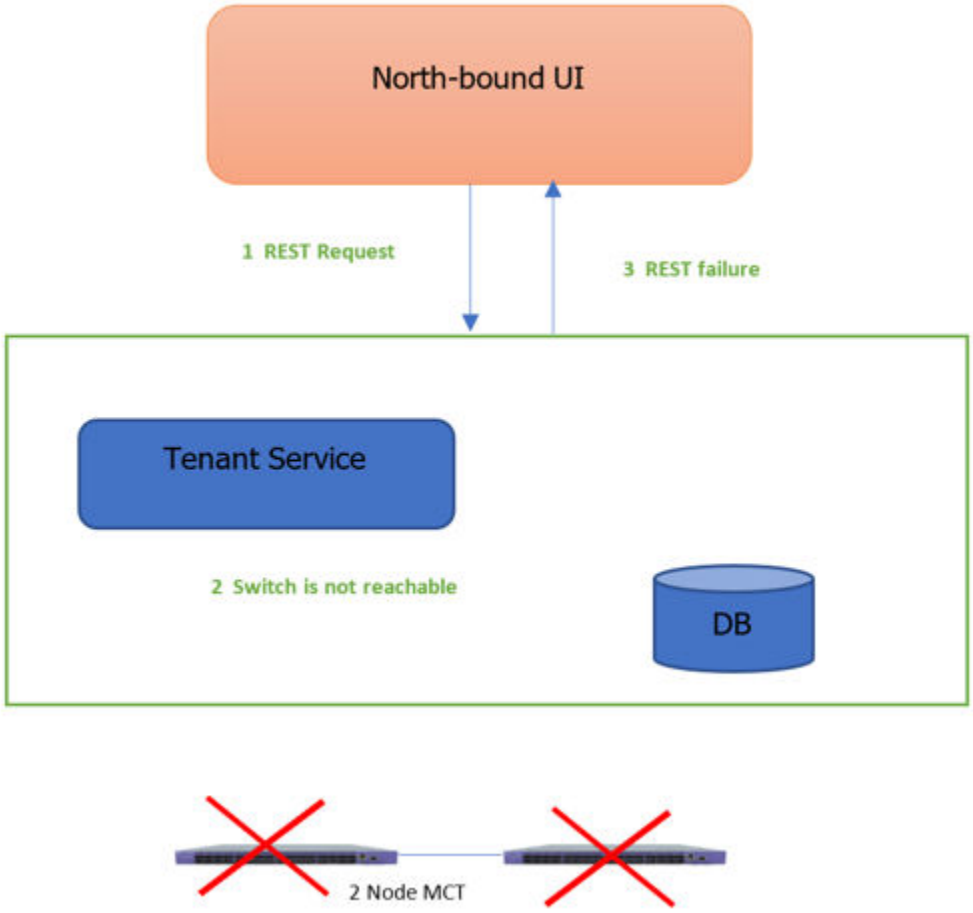

Figure 7: Rollback for failure of both nodes

Failure during De-configuration

Rollback does not occur when a REST operation fails during a de-configuration request. The status of configuration items that were not rolled back changes to "delete-pending." You must manually verify and address the status of such items.

XCO High Availability Failover Scenarios

XCO high availability provides for uninterrupted service in several different scenarios.

For information about deploying XCO for high availability, see the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

SLX Device Failure

When an SLX device fails, the SLX-OS and the XCO services running on TPVM go down for the failed node. The time it takes for failover to the standby node varies depending on whether the K3s agent node is actively running the XCO services. The following image depicts a scenario in which one SLX device fails.

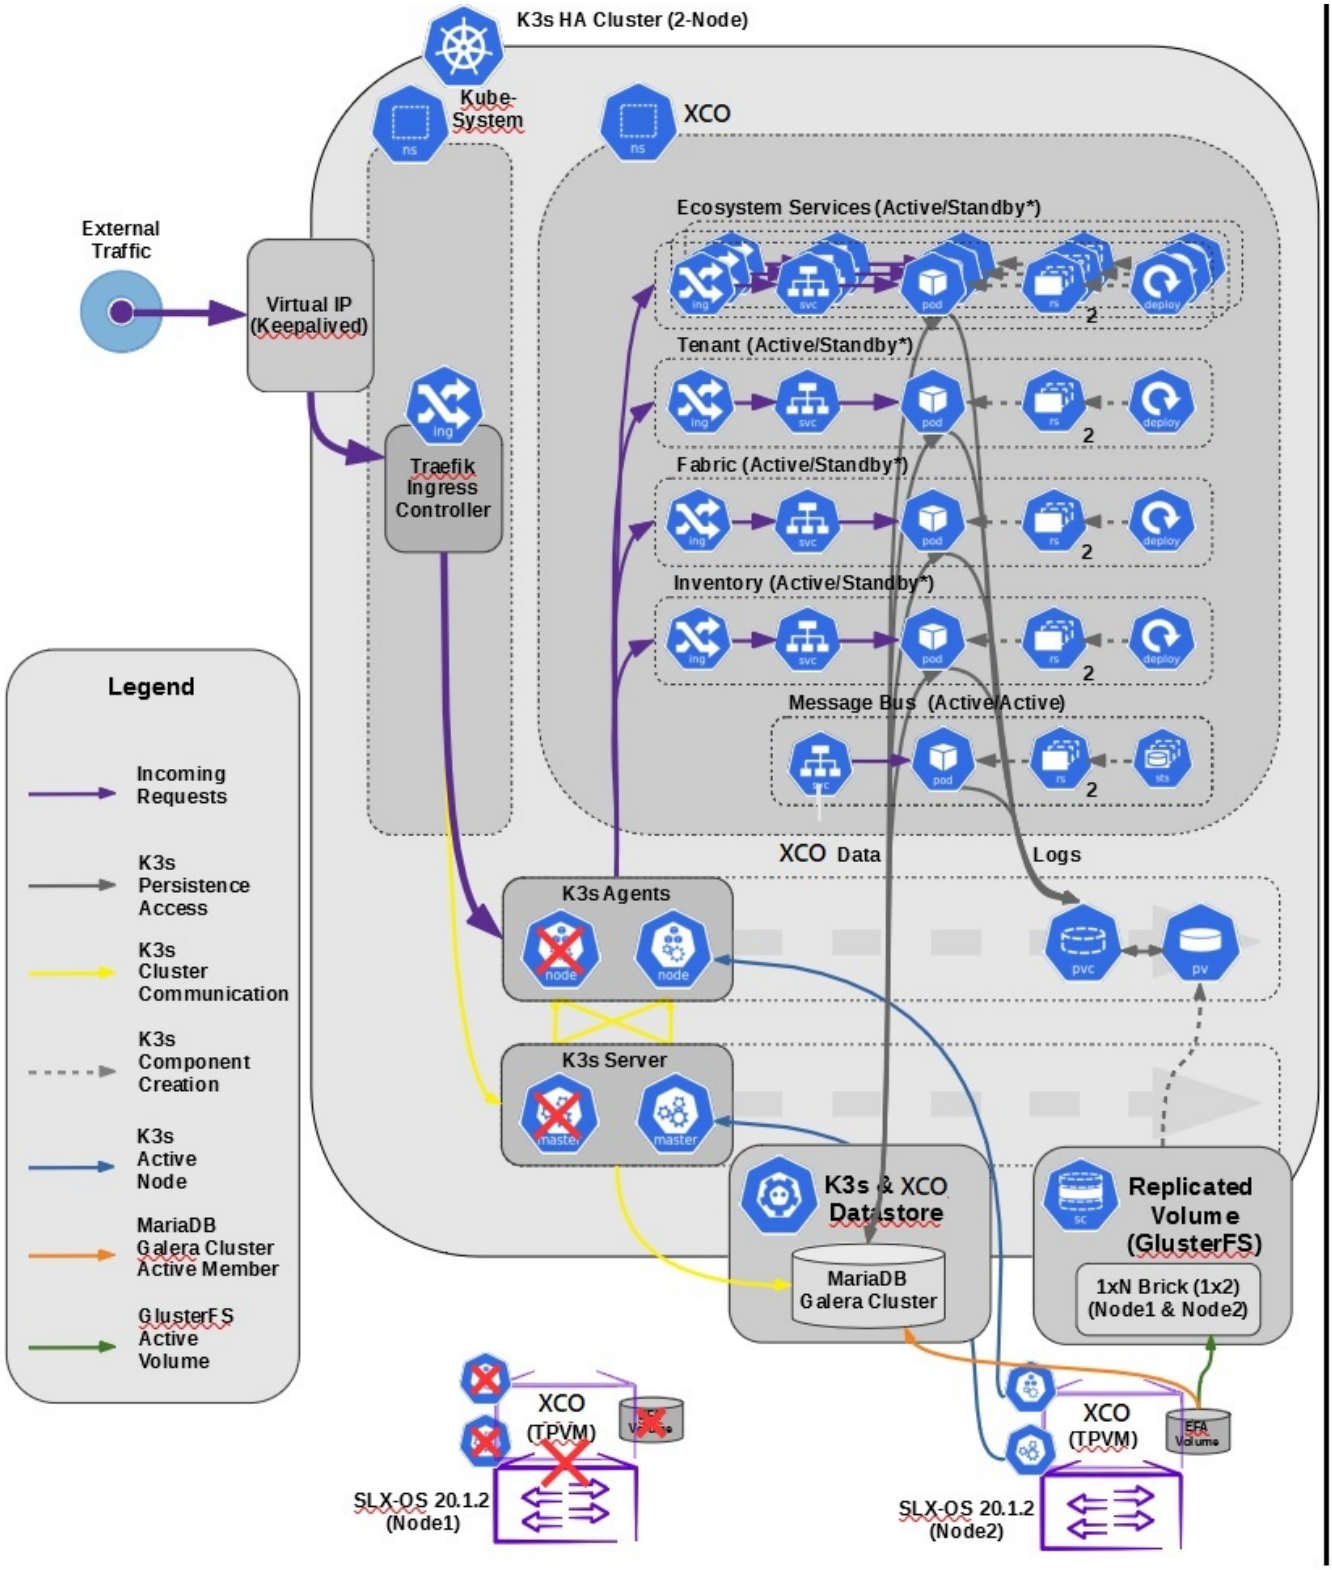

Figure 8: SLX device failure in a two-node cluster

SLX Device Failure on the Active K3s Agent Node

When the K3s agent node is actively running XCO services on a node that fails, K3s initiates failover and starts the XCO services on the standby node. Failover is complete when XCO services are running on the newly active K3s agent node (node 2).

Because the GlusterFS replicated volume remains available during failover, the K3s cluster data store and the XCO data store remain operational.

When the failed node is again operational, it becomes the standby node. The K3s agent node continues to run XCO services from node 2. When both nodes are up and K3s is running, all services fetch the latest data from devices to ensure that XCO has the latest configurations.

SLX Device Failure on the Standby K3s Agent Node

When the K3s agent node is the standby and is not running XCO services, no failover actions occur if this node fails. XCO services continue to run on the active node without interruption.

TPVM Failure

The TPVM failure scenario is similar to that of the SLX device failure scenario. The only difference is that SLX-OS continues to operate.

Two-node Failure

In the unlikely event that both nodes in the cluster fail at the same time (for reasons such as a power failure or the simultaneous reboot of SLX devices), XCO has built-in recovery functionality. If the cluster is not automatically recovered within 10 minutes of power being restored or within 10 minutes of the TPVM being rebooted, then you can manually recover the cluster.

Multiple Management IP Networks

Overview

The Multiple Management IP (MMIP) Networks feature offers the following support:

- Supports single node and multi-node deployments
- Supports TPVM deployments, server-based deployments, and VM-based deployments
- Supports the configuration of additional management IP networks and routes during XCO installation
- Supports adding and viewing management networks and routes after XCO installation
- Supports deleting management networks and routes after XCO installation
- Supports the migration of the multiple network configuration during the following XCO upgrade scenarios: single node to multi-node and multi-node to multi-node

- Supports up to 6 networks
- Supports the RMA, backup, restore, and upgrade functions

Note
If you do not need multiple management networks, simply reply "no" when prompted during XCO installation or upgrade. For instructions, see the installation and upgrade topics in the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

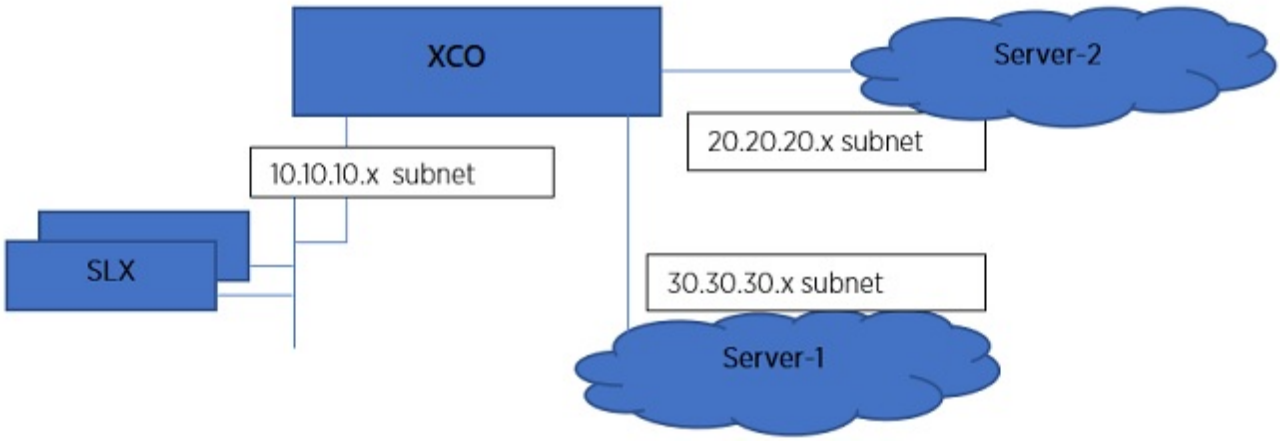

Figure 9: MMIP architecture

Assumptions

- In a multi-node deployment, the sub-interface with the VLAN is created under the same NIC as the VIP destination. In a single-node deployment, the sub-interface is created under the NIC that you specified as the host IP installation (if there are multiple NICs). Creating sub-interfaces on different NICs of the server is not supported.
- XCO does not validate connectivity to the newer IP subnets. You are responsible for ensuring reachability.
- Changing IP subnets or IP routes is not supported. To make changes to a management network or network route, you must delete the network or route and then create a new one.
- You can expect about 20 to 30 seconds of downtime when adding or deleting management networks.
- In a multi-node deployment, both nodes have to be up and available during add and delete operations (because sub-interface creation and keepalived changes are unique to the node). Because these are infrequent operations, you should verify that both the nodes are up and in READY state before beginning add or delete operations.

Add and Delete Management Routes

In a multi-node deployment, you can add, delete, and show management routes for Multiple Management IP (MMIP) networks.

About This Task

Typically, the create and delete operations do not cause a high-availability failover. The route is instantiated on the active node of the cluster. If failover does occur from node 1 to node 2, then keepalived ensures that the route transitions from node 1 to node 2.

Procedure

1. To add a management route, run the following command.

```
$ efa mgmt route create --src <mmip-vip> --to <dest-cidr> --via <next-hop-ip>
```

   a. Run the following command to list both IPv4 and IPv6 routes:

```
$ efa mgmt route show
+------------+------------+------------+
| Route-Src  |  Route-To  | Route-Via  |
+------------+------------+------------+
| 2000::1    | ffee::/64  | 2000::2    |
+------------+------------+------------+
| 2000::1    | 4000: :/64 | 2000::3    |
+------------+------------+------------+
| 10.10.10.1 | 1.1.1.0/24 | 10.10.10.2 |
+------------+------------+------------+
```

   b. Run the following command to list IPv4 routes:

```
$ efa mgmt route show ipv4
+------------+---------------+------------+
| Route-Src  |   Route-To    | Route-Via  |
+------------+---------------+------------+
| 10.10.10.1 | 11.11.11.0/24 | 10.10.10.2 |
+------------+---------------+------------+
| 10.10.10.1 | 12.12.12.0/24 | 10.10.10.2 |
+------------+---------------+------------+
```

   c. Run the following command to list IPv6 routes:

```
$ efa mgmt route show ipv6
+-----------+-----------+-----------+
| Route-Src | Route-To  | Route-Via |
+-----------+-----------+-----------+
| 3000::1   | 4000::/64 | 3000::2   |
+-----------+-----------+-----------+
| 3000::1   | 5000::/64 | 3000::2   |
+-----------+-----------+-----------+
```

   If a route with the same destination exists, the operation fails. This operation updates the keepalived configuration file on both nodes of the high-availability cluster.

2. To delete a management route, run the following command.

```
$ efa mgmt route delete --src <mmip-vip> --to <dest-cidr> --via <next-hop-ip>
```

   If a route matching the three parameters does not existing, the operation fails. If a matching route is found, the keepalived configuration file is updated and reloaded on both nodes of the high-availability cluster.

3. To generate a list of all management routes, run the following command.

```
$ efa mgmt route show
+-------------+------------------+-------------+
|  Route-Src  |     Route-To     |  Route-Via  |
```

```
+-------------+------------------+-------------+
| 10.21.30.40 | 192.168.100.0/24 | 10.21.30.41 |
+-------------+------------------+-------------+
```

Configuration Supporting Multiple Management IP Networks

This work flow highlights the changes that occur in your system when you configure Multiple Management IP (MMIP) networks.

Day 0 and Installation Configuration

- In a multi-node deployment, the VIP (virtual IP address) that you enter during installation is the same as for a non-MMIP deployment. This VIP is distinguished from those added during MMIP network operations and cannot be deleted.
- During installation, you are prompted to create additional MMIP networks and routes.
- Keepalived, ingress, and interface changes are performed on both nodes of a multi-node deployment.
- Configuration is persisted for RMA purposes, so that the Supportsave function has data for debugging issues.

For step-by-step instructions for configuring MMIP during installation or upgrade, see the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

Day 1 to Day *n* Configuration

- You can use the XCO CLI or REST APIs to add and delete management routes and IP address and VLAN combinations.
- Keepalived, ingress, and interface changes are performed on both nodes of a multi-node deployment.
- Configuration is persisted for RMA purposes, so that the Supportsave function has data for debugging issues.
- The backup and restore process also restores the previous configuration of the sub-interfaces.

Add and Delete Management Sub Interfaces

You can use the XCO CLI to add and delete management sub interfaces.

About This Task

Follow this procedure to add or delete a management sub interface.

Procedure

1. Run the following command to add a management sub interface:

```
efa mgmt subinterface create [ --name sub | --vlan-id vlan-id | --ip-addr ip-addr --
ipv6-address ipv6-addr ]
```

   a. Run the following command to show IPv6 routes:

```
#efa mgmt subinterface show
+------+-----------------+------+---------------+-------------+
| Name | Parent Interface| Vlan |   IP Subnet   | IPv6 Subnet |
+------+-----------------+------+---------------+-------------|
```

```
| sub1 | ens160          | 100  | 10.10.10.1/24 |             |
+------+-----------------+------+---------------+-------------+
| sub2 | ens160          | 200  | 11.11.11.1/24 | 2000::1/64  |
+------+-----------------+------+---------------+-------------+
| sub3 | ens160          | 300  | 13.13.13.1/24 |             |
+------+-----------------+------+---------------+-------------+
| sub4 | ens160          | 400  | 14.14.14.1/24 | 3000::1/64  |
+------+-----------------+------+---------------+-------------+
```

- If a management network with the same name exists, the operation fails.
- The changes made by this operation span three different components:
  - Sub interface creation under the physical NIC
  - Keepalived configuration changes (for high-availability deployments)
  - Ingress controller changes
- If any operation to the component fails, it is marked as a failed operation and the configurations return to the previous state.

2. Run the following commands to delete a management sub interface:

```
efa mgmt subinterface delete --name <name>
```

If a management network with the name exists, it is deleted. Otherwise, the correct response is provided in the command output.

Example

```
$ efa mgmt subinterface?
Management subinterface commands

Usage:
  efa mgmt subinterface [command]

Available Commands:
  create        Create sub-interface (sub-interface)
  delete        Delete sub-interface (sub-interface)
  show          List of sub-interfaces (sub-interfaces)

Use "efa mgmt subinterface [command] --help" for more information about a command.
$ efa mgmt subinterface create -h
Create management subinterface (sub-interface)
Usage:
  efa mgmt subinterface create [flags]
Flags:
      --name string         Name of the sub-interface
      --vlan-id int         VLAN Id of sub-interface
      --ip-address string   IP Address of sub-interface including subnet mask.
$ efa subinterface delete -h
Delete management subinterface (sub-interface)
Usage:
  efa mgmt subinterface delete [flags]
Flags:
      --name string   Name of the sub-interface
$ efa mgmt subinterface show -h
List of management sub-interfaces (sub-interfaces)

Usage:
  efa mgmt subinterface show [flags]
```

```
 Flags:
       --name string   Name of the sub-interface

$ efa mgmt subinterface create --name server1 --vlan-id 20 --ip-address
          20.20.20.2/24
Subinterface server1 created successfully

$ efa mgmt subinterface delete --name server1
Subinterface server1 deleted successfully

$ efa mgmt subinterface show
+------+------------------+------+----------------+-------------+
| Name | Parent Interface | Vlan |   IP Subnet    | IPv6 Subnet |
+------+------------------+------+----------------+-------------+
| sub1 | ens160           | 100  |  10.10.10.1/24 |             |
+------+------------------+------+----------------+-------------+
| sub2 | ens160           | 200  |  11.11.11.1/24 | 2000::1/64  |
+------+------------------+------+----------------+-------------+
| sub3 | ens160           | 300  |  13.13.13.1/24 |             |
+------+------------------+------+----------------+-------------+
| sub4 | ens160           | 400  |  14.14.14.1/24 | 3000::1/64  |
+------+------------------+------+----------------+-------------+
Management Subinterfaces Details

$ efa mgmt subinterface show --name server1
+---------------+------------------+------+---------------+
| Sub-Interface | Parent Interface | Vlan |   IP Subnet   |
+---------------+------------------+------+---------------+
| server1       | eth0             | 20   | 20.20.20.2/24 |
+---------------+------------------+------+---------------+
Management Subinterface Details
```

Configure Static IP Addresses for Management Sub interfaces

You can use the XCO CLI to add, delete, and show the static IP addresses for management sub interfaces.

About This Task

Follow this procedure to configure static IP address for management sub interface.

Procedure

1. To add static IP addresses to a specified sub interface, run the following command.

   ```
   efa mgmt subinterface staticip add [ --subinterface sub | --ip1 ip1 | --ip2 ip2]
   ```

   Here ip1 and ip2 can be either IPv4 or IPv6 address. All other assumptions and limitations for IPv6 static IPs are same as IPv4 static IPs. You can also configure a mix of IPv4 and IPv6 address for ip1 and ip2.
2. To remove static IP addresses from a specified sub interface, run the following command.

   ```
   efa mgmt subinterface staticip remove
   --subinterface <int-id>
   ```
3. To show all sub interfaces and the IP addresses that are attached to them, run the following command.

   ```
   efa mgmt subinterface staticip show
   ```

Change the Default Gateway of a TPVM

You can change the default gateway of a TPVM.

About This Task

Follow this procedure to change the default gateway of a TPVM. This procedure has no impact on the functioning of XCO high-availability mode.

- The gateway IP address must be in the same subnet as one of the subinterfaces that are created in XCO.
- To configure a default gateway for the subinterface subnet, use only static IPv4 addresses.
- Perform this procedure on both nodes where XCO is deployed, to avoid a loss of XCO functionality.
- Perform this procedure with caution. XCO and SLX-OS do not validate reachability of the gateway during this operation.
- Vital services, such as DNS, NTP, and LDAP, must be reachable from the new gateway.

Procedure

1. Add static IP addresses to the sub interface.

   You can assign a maximum of one pair of static IP addresses. Only one sub interface at a time can have static IP addresses.

   ```
   efa mgmt subinterface staticip --name <sub-int-name> --ip1 <ip-addr>
   --ip2 <ip-addr>
   ```

2. Change the gateway of the standby TPVM.

   ```
   efa inventory device execute-cli --ip <standby-slx-ipaddr>
   --command "tpvm TPVM,interface management ip <standby-tpvm-ipaddr>
   gw <new-gateway-ipaddr>" --config
   ```

3. Change the gateway of the active TPVM.

   ```
   efa inventory device execute-cli --ip <active-slx-ipaddr>
   --command "tpvm TPVM,interface management ip <active-tpvm-ipaddr>
   gw <new-gateway-ipaddr>" --config
   ```

Northbound IPv6 support

EFA 3.1.0 supports IPv6 via an additional virtual IP that is assigned during EFA installation or upgrade. Clients can make REST API calls to XCO via either an IPv4 or an IPv6 address. An IPv6 address is optional, but an IPv4 address is mandatory.

You can configure the API gateway with IPv4 and IPv6 addresses.

Note

- Dual mode or IPv4 only mode is supported. TPVM does not support assigning an IPv6 address via the SLX CLI.
- (Optional) Upgrades from earlier releases or from single-node to multi-node and fresh install require an IPv6 address configuration.
- Single-node installation supports IP address change for both IPv4 and IPv6 addresses.
- API gateway supports IPv6 virtual IP address for multi-node and single-node installation.

All the northbound CLIs, which support dual IP mode, validate IP address according to the derived mode.

For information about commands and supported parameters to configure IPv6, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Verify the Upgrade

1. On the SLX device, run the **`show efa status`** command to see details of the installation and the state of services.
2. From the XCO command line, run the **`efactl status`** command to see the status of nodes, pods, and services.
3. Run the **`efa status`** command for concise status information.

Note

For information about upgrading XCO, see the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

Verify the IPv6 Address

During fresh installation or upgrade, ensure that you have specified a valid IPv6 address.

1. When prompted, select **Yes** to assign an IPv6 address as a virtual IP for cluster.
2. Specify the virtual IPv6 address. The system validates the IPv6 address to ensure that it is not in use. If it is already in use, change the IPv6 address.

Configure DNS Nameserver Access

A well configured DNS server during XCO deployment enables XCO services access to a host DNS nameserver.

About This Task

Follow this procedure to enable XCO service access to a host DNS nameserver. Use the script available in the `/apps/efa/` directory on a TPVM and in the `/opt/efa` directory on the server.

Note

- Ensure that you are a root user or have sudo privileges.
- Ensure that the DNS nameserver is valid.
- In a multi-node deployment, ensure that you update the DNS nameserver on both nodes.

Procedure

1. To enable XCO services access to a host DNS nameserver, run the following command:

```
sudo <location of the script>/update-dns.sh --dns-action allow
```

2. To disable XCO services access to a host DNS nameserver, run the following command:

```
sudo <location of the script>/update-dns.sh --dns-action disallow
```

Change Password of efainternal User

A new user created during installation is named as efainternal user.

About This Task

The installation or upgrade of XCO creates a new user on the host with a random password. The name of this new user is efainternal user. Prior to XCO 3.2.0 (EFA 3.1.0 or earlier), changing the password of an efainternal user impacts the functionality of EFA.

Procedure

1. To update the password of 'efainternal' user in XCO, run the following script:

   The script is available in the `/apps/efa/` in TPVM and `/opt/efa` directory on a server.

```
extreme@tpvm:~$ sudo bash /apps/efa/update-password.sh --help

/opt/efa/update-password.sh Usage:

--help - show this message

--username <user_name>, name of the user

--password <password>, - OPTIONAL, password for the user

--random-password, -OPTIONAL, sets a random password to the user
```

```
--update-reference-only, -OPTIONAL, updates reference without any password change,
this is applicable for 'efainternal' user
```

2. To assign a random password, skip the password parameters.
3. To manually update a password on the host of all the nodes of XCO, run the following script with the `update-reference-only` parameter:

```
user@ubuntu:/opt $ sudo bash /opt/efa/update-password.sh --username efainternal

Password:

Saving EFA user information for this node

Password update is successful.
```

Accessing Supportsave without Login

In XCO 3.2.0, without logging into the XCO, the monitoring service uses a system generated token to fetch data (supportsave file) from the system. The token does not have an expiry date.

A new `ConfigReader` role is added to the token with the permission to get running configuration and debug-locks. The `ConfigReader` role is an internal role like a 'ServiceAdmin' and cannot be assigned to any user.

Linux Exit Codes

From XCO 3.2.0, errors found while running a command will return a Linux exit code of 1.

Linux Error Exit Code

The following example shows that any device failure will return 1 (error).

```
$ efa inventory device interface set-admin-state --ip 10.139.44.175-177 --if-type eth --
if-name 0/1 --state up
+---------------+-----+------+-----------+--------+---------+-------------------------+
|   DeviceIP    | ID  | Name | Interface | Admin  | Result  |         Reason          |
|               |     |      | Type      } Status |         |                         |
+---------------+-----+------+-----------+--------+---------+-------------------------+
| 10.139.44.176 |     |      |           |        | Failed  | Device does not exist   |
|               |     |      |           }        |         | with IP: 10.139.44.176 |
+---------------+-----+------+-----------+--------+---------+-------------------------+
| 10.139.44.177 |     |      |           |        | Failed  | Device does not exist   |
|               |     |      |           }        |         | with IP: 10.139.44.177 |
+---------------+-----+------+-----------+--------+---------+-------------------------+
| 10.139.44.175 | 297 | 0/1  | ethernet  | up     | Success |                         |
+---------------+-----+------+-----------+--------+---------+-------------------------+
Interface Details
--- Time Elapsed: 23.384583544s ---
```

```
$ echo $?
1
```

Note

The last line shows 0 instead of 1 even when at least one "Failed" result is reported.

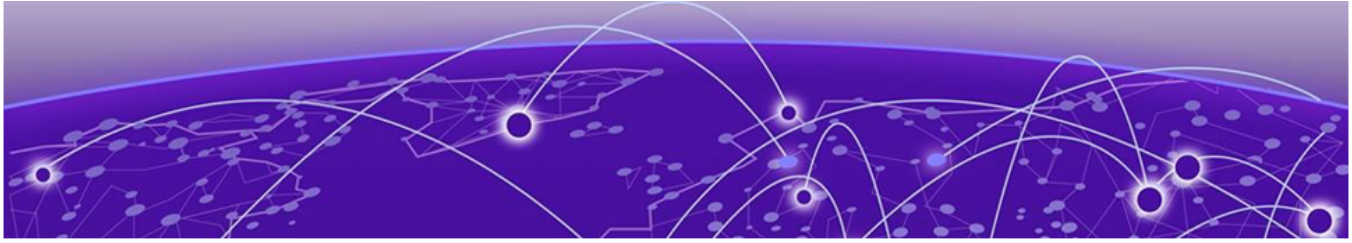

Fabric Infrastructure Provisioning

Fabric Service Overview

Fabric Service is responsible for automating the Fabric BGP underlay and EVPN overlay. By default, the EVPN overlay is enabled but can be disabled before provisioning if desired. Fabric Service exposes the CLI and REST API to clients for automating the fabric underlay and overlay configuration.

Fabric Service features include:

- Small Data Center Topology (small data center support)
- Support for 3- and 5-stage Clos fabrics
- Support for MCT configuration
- Support for Eco-System Integration; Openstack, VMWare vCenter, Microsoft Hyper-V, and SCVMM

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. Fabric Service is deployed along with Inventory Service and Tenant Service.

Note

You cannot perform fabric and tenant operations when manual DRC is in progress.

IP Fabric and Clos Orchestration Overview

A fabric is a logical container for holding a group of devices. Here it denotes a collection of devices that are connected in a fabric topology and on which you can configure underlay and overlay.

Fabric service provides following features:

- 3-stage Clos automation
- 5-stage Clos automation
- Small Data Center automation
- Multi-Fabric automation
- Fabric topology view
- Fabric validation, error reporting, and recovery
- Single-homed leaf or multi-homed (MCT) leaf

Fabric CLIs and REST APIs provide the following:

- Mechanism to create a fabric composed of multiple DC points of delivery (PoDs).
- Mechanism to configure fabric settings. Fabric settings are collections of settings that control the various parameters of the fabric being managed, for example, Layer 2 and Layer 3 MTU, and BGP maximum paths.

  For more information about the commands, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*

  .
- Mechanism to fetch per-device errors occurring during fabric configuration, for which you can take corrective or remedial actions.

Errors occurring on the device during fabric creation are tagged against the devices and can be retrieved from the CLI and REST APIs for use in taking corrective or remedial actions.

SLX Device Prerequisites for Fabric Service

The following items are required before you configure your fabric.

- Management IP addresses must be configured on all devices.
- SLX devices must have the appropriate firmware version. For more information, see the list of supported platforms in the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.
- SLX 9850: Fabric links must be enabled manually, through `no shut`.
- SLX 9540: The appropriate TCAM profile must be set and the device rebooted.

```
device# conf
Entering configuration mode terminal
device(config)# hardware
device(config-hardware)# profile tcam vxlan-ext
%Warning: To activate the new profile config, run 'copy running-config startup-config'
```

```
followed by 'reload system'.
device(config-hardware)#
```

- Refer to the release-specific *Extreme SLX-OS Management Configuration Guide* for configuration steps for each platform.

Configure Hardware Profile to Limit IPv6 Prefix to 64

You can configure a hardware profile to limit the maximum length of IPv6 prefix to 64. The hardware profile configuration lets you increase the scale of IPv6 prefix installed on the routing hardware.

About This Task

Follow this procedure to configure a hardware profile.

Note

- The hardware profile configuration is applicable only to Extreme 8520, 8720, 8820, 9150, and 9250 hardware on the SLX firmware version 20.2.1 and above.
- You can configure CMD on the deployed fabric follow by warning message to reboot the system
- The default value (of `maximum-ipv6-prefix-length-64 string`) is "" empty string. You can configure **Yes** or **No**.
- If you configure device settings at fabric level and then add a new device to the fabric, the configured device settings at fabric level will not be applicable to the newly added device. You need to configure the device settings to the newly added device as shown in the following example:

```
- efa inventory device setting update --maximum-ipv6-prefix-length-64 Yes --
urpf Yes --fabric fs
        - Adding a new device-3 to the fabric (configured settings --urpf yes
at fabric level won't derived and effect to device-3)
        - efa inventory device setting update --maximum-ipv6-prefix-length-64
Yes --urpf Yes -ip device-3 (user again configure the value)
```

Procedure

1. Run the following command:

```
efa inventory device setting update [flags]

      --ip string                                        Specifies a comma-separated
range of device IP addresses. For example: 1.1.1.1-3,1.1.1.2,2.2.2.2
      --fabric string                                    Specify the name of the fabric
      --maint-mode-enable-on-reboot string               Enter Yes to configure maintenance
mode enable on reboot and No to de-configure
      --maint-mode-enable string                         Enter Yes to configure maintenance
mode enable and No to de-configure
      --maint-mode-convergence-time string               Maximum time in seconds that
maintenance mode is allowed to complete operations, valid values 100-500 and 0 to
de-configure
      --mct-bring-up-delay string                        Delay, in seconds, waited before MCT
cluster bring-up, valid values 10-600 and 0 to de-configure
      --health-check-enable string                       Enter Yes to enable health check and
No to disable health check
      --health-check-interval string                     Health check interval in seconds/
minutes, valid values for Fabric device 6m-24h, valid values for NPB device 30s-24h
```

```
Example. 30s or 99m or 1h20m or 20m, default 6m for Fabric device, 30s for NPB device
      --health-check-heartbeat-miss-threshold string   Health check's heartbeat miss
threshold value, valid value range in between 2-5, default 2
      --config-backup-periodic-enable string           Enter Yes to enable periodic config
backup and No to disable periodic config backup
      --config-backup-interval string                  Config Backup interval in minutes,
valid values 3m-30h Example. 99m or 1h20m or 20m , default 24h
      --number-of-config-backups string                Config Backup Count, valid values
2-20, default 4
      --prefix-independent-convergence string          Enter Yes to enable BGP PIC and No
to de-configure
      --prefix-independent-convergence-static string   Enter Yes to enable Static PIC and
No to de-configure
      --maximum-load-sharing-paths string              Config route load-sharing maximum
paths, valid values 8,16,32,64, default 64 paths
      --maximum-ipv6-prefix-length-64 string           Enter Yes to configure the
maximum route prefix length of 64, valid values Yes/No, default "". This configuration
is applicable for SLX-9150, SLX-9250, Extreme 8720 and Extreme 8520 hardware
      --urpf string                                    Enter Yes to configure the
unicast reverse path forwarding, valid values Yes/No, default "". This configuration
is applicable for SLX-9150, SLX-9250, Extreme 8720 and Extreme 8520 hardware
```

2. Complete the following configuration on SLX device:

```
Rack1-Device1# show running-config hardware
hardware
   profile route enable ipv6-max-prefix-64 urpf
!
```

Example

The following example configures a hardware profile that limit the maximum length of IPv6 prefix to 64:

```
$ efa inventory device setting update --maximum-ipv6-prefix-length-64 Yes --urpf Yes --ip
10.20.48.93

+-------------+--------------------------------+---------+-------+-------+
| IP ADDRESS  |              NAME              | STATUS  | VALUE | ERROR |
+-------------+--------------------------------+---------+-------+-------+
| 10.20.48.93 | Maximum Ipv6 Prefix Length 64 | Success | Yes   |       |
+             +--------------------------------+---------+-------+-------+
|             | Urpf                           | Success | Yes   |       |
+-------------+--------------------------------+---------+-------+-------+

Warning: Maximum Ipv6 Prefix Length 64 configuration will not take effect until reloaded.

Execute the CLI to reload : efa inventory device reload --ip 10.20.48.93

Warning: Urpf configuration will not take effect until reloaded.

Execute the CLI to reload : efa inventory device reload --ip 10.20.48.93

--- Time Elapsed: 14.1348949s ---


$ efa inventory device setting show --ip 10.20.48.93

+---------------------------------+-------+
|              NAME               | VALUE |
+---------------------------------+-------+
| Maintenance Mode Enable On      | No    |
| Reboot                          |       |
+---------------------------------+-------+
| Maintenance Mode Enable         | No    |
+---------------------------------+-------+
```

```
| Maintenance Convergence Time   |       |
+--------------------------------+-------+
| MCT Bring-up Delay             |       |
+--------------------------------+-------+
| Health Check Enabled           | No    |
+--------------------------------+-------+
| Health Check Interval          | 6m    |
+--------------------------------+-------+
| Health Check Heartbeat Miss    | 2     |
| Threshold                      |       |
+--------------------------------+-------+
| Periodic Backup Enabled        | Yes   |
+--------------------------------+-------+
| Config Backup Interval         | 24h   |
+--------------------------------+-------+
| Config Backup Count            | 4     |
+--------------------------------+-------+
| Prefix Independent Convergence | No    |
+--------------------------------+-------+
| Static Prefix Independent      | No    |
| Convergence                    |       |
+--------------------------------+-------+
| Maximum Load Sharing Paths     | 64    |
+--------------------------------+-------+
| Maximum Ipv6 Prefix Length 64  | Yes   |
+--------------------------------+-------+
| Urpf                           | Yes   |
+--------------------------------+-------+
--- Time Elapsed: 56.1149ms ---
```

Clos Overview

XCO offers unique flexibility in supporting 3- and 5-stage Fabric Clos topologies based on a BGP underlay with a BGP or EVPN overlay.

Tenant Network onboarding services are supported on both topologies, allowing you to create connectivity for devices connected to the fabric, such as compute (servers), storage, and connectivity to external routers or gateways.

3-Stage Clos

3-stage Clos consists of an ingress leaf layer, a middle spine layer, and an egress leaf layer. Servers are connected to leaf devices and leaf devices are connected to all spines. No leaf devices are connected to other leaf devices, nor are spines connected to spines. Data enters at an ingress leaf, is routed through a spine to an egress leaf, and then out of the network to the next server in the path. In this topology, servers are always 3 hops (leaf, spine, leaf) away from another server.

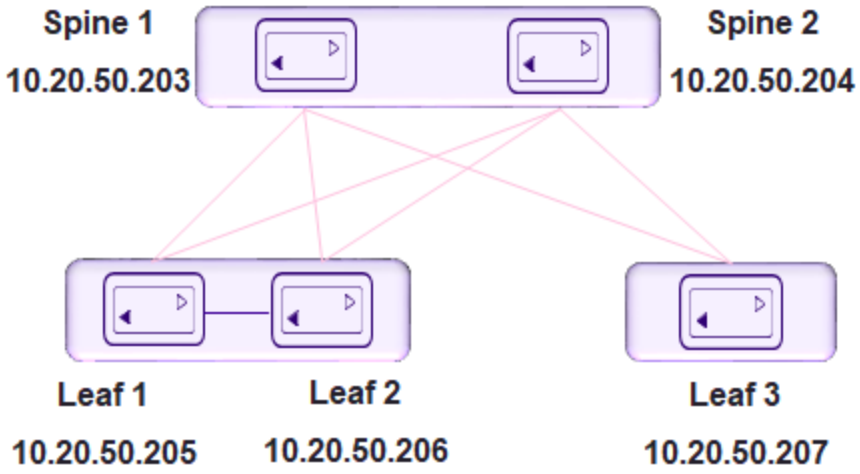

5-Stage Clos

5-stage Clos is a 3-stage topology that is divided into clusters and on which a Super-spine layer is added. All links between leaf and spine must be connected. Spine are not be interconnected. Similarly, all the links between the spine and Super-spine must be connected.

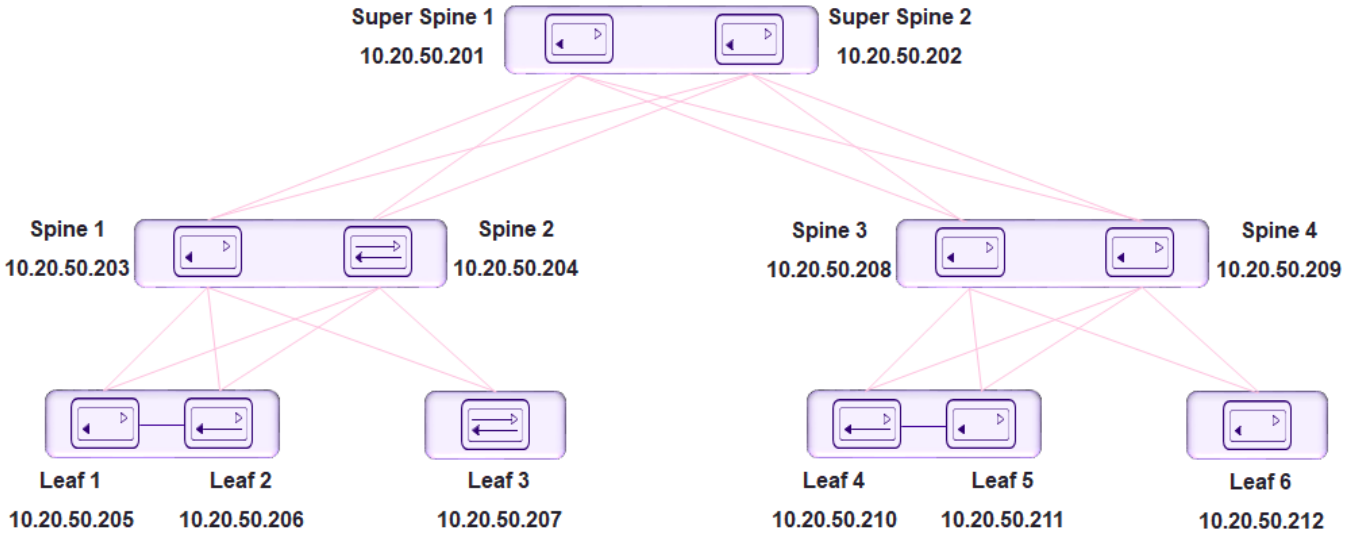

Configure a 3-Stage Clos Fabric

The 3-stage topology has 2 layers of devices: leaf and spine. All links between leaf and spine must be connected. Spine nodes are not interconnected.

About This Task

Tip
If any devices in a fabric are in "admin-down" state, use of the following commands in that same fabric will not add or delete devices in the fabric: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

Procedure

1. Create the fabric.

   ```
   efa fabric create
   ```

2. Add a device to the fabric.

   ```
   efa fabric device add
   ```

   A device must be registered with the Inventory Service before you can add it to a fabric. However, if you provide a user name and password when you run the command, then the devices are automatically registered with the Inventory Service. See the examples at the end of this procedure.

   You can add multiple devices by using the **`efa fabric device add-bulk`** command.

   

   Tip
   To validate fabric port-link status, complete the following operations before running the **`efa fabric device add-bulk`** command:

   a. Run the **`efa inventory device register -ip <list of device-ips>`** command.
   b. Run the **`efa inventory device interface list -ip <device-ip>`** command.
      i. Verify port link status (up or down) in Admin Status and Oper Status fields.
      ii. Confirm they are as expected.
      iii. If not, manually check for physical cabling and fix any issues. Continue with the `efa fabric device add-bulk` operation.

3. Configure the fabric.

   ```
   efa fabric configure
   ```

   Topology validation occurs during the addition of a device and during fabric configuration. The following validations are performed.

   - Leaf nodes must connect to all the spine nodes.
   - A spine node must connect to all the leaf nodes.
   - A border leaf node connects to all the spine nodes.

- A spine node connects to all the border leaf .
- No more than two leaf nodes connect to each other.
- No more than two border leaf nodes connect to each other.
- Border leaf node and leaf node are not connected to each other.
- Spine nodes are not connected to each other.
- Super-spine nodes are not connected to each other.
- A leaf node marked as "multi-homed" must have an MCT neighbor.
- A leaf node marked as "single-homed" is not connected to other leaf nodes.
- A border beaf node marked as "multi-homed" must have an MCT neighbor.
- A border leaf node marked as "single-homed" is not connected to other border leaf nodes.
- Device role (such as leaf, border-leaf, spine, or super-spine) is validated for a given device platform type (for example, SLX 9840 cannot be added as a leaf).

Tip
The validation process reports any errors as a response to the **`efa fabric device add`** or **`efa fabric configure`** operations. You can use the **`efa fabric error show`** command to export these errors to a CSV file.

Note
You cannot change fabric settings after you add devices to the fabric, with the following exceptions: `--md5-password-enable`, `--md5-password`, `--bgp-dynamic-peer-listen-limit`, and `--single-rack-deployment settings`.

Example
This example creates the fabric.

```
efa fabric create --name stage3
```

Example
This example adds multiple devices to the fabric.

```
efa fabric device add-bulk --leaf 10.20.50.205,10.20.50.206,10.20.50.207
--spine 10.20.50.203,10.20.50.204 --name stage3 --username admin
--password password
```

Example
This example configures the fabric.

```
efa fabric configure --name stage3
```

Configure a 5-Stage Clos Fabric

The 5-stage topology has three layers of devices: leaf, spine, and super-spine.

About This Task
You can build a 5-stage Clos from top to bottom or bottom to top. The following example builds from top to bottom.

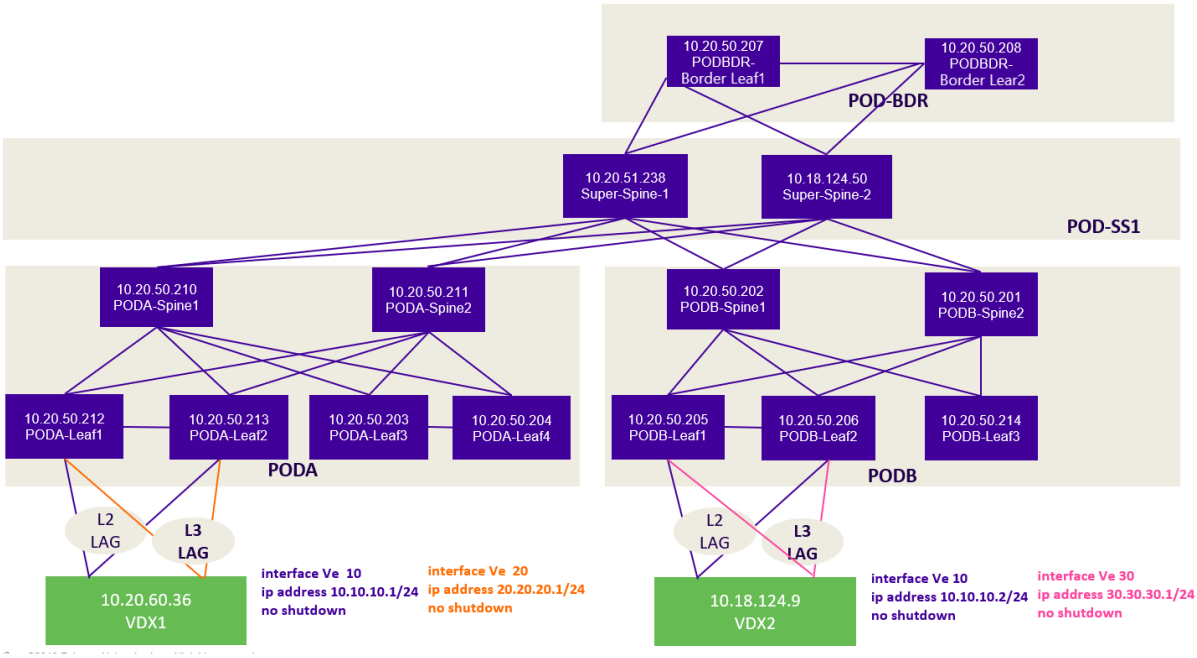

Figure 10: 5-Stage Clos fabric topology

Tip

If any devices in a fabric are in "admin-down" state, use of the following commands in that same fabric will not add or delete devices in the fabric: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

Procedure

1. Create the fabric.

   ```
   efa fabric create
   ```

2. Add a device to the fabric.

   ```
   efa fabric device add
   ```

   A device must be registered with the Inventory Service before you can add it to a fabric. However, if you provide a user name and password when you run the command, then the devices are automatically registered with the Inventory Service. See the examples at the end of this procedure.

   You can add multiple devices by using the **`efa fabric device add-bulk`** command. If you choose to add multiple devices in bulk, ensure you perform the following operations first:

   - Run the **`efa inventory device register --ip <list-of-device-ips>`** command.
   - Run the **`efa inventory device interface list --ip <device-ip>`** command. In the output of the command, verify that the states of the port links are as you expected (in the Admin Status and Oper Status fields). If not, manually check the physical cabling and fix any issues. Then continue with the **`efa fabric device add-bulk`** operation.

3. Configure the fabric.

```
efa fabric configure
```

Topology validation occurs during the addition of a device and during fabric configuration. The following validations are performed:

- Leaf nodes must connect to all the spine nodes.
- A spine node must connect to all the leaf nodes.
- A border leaf node connects to all the spine nodes.
- A spine node connects to all the border leaf nodes.
- No more than two leaf nodes connect to each other.
- No more than two border leaf nodes connect to each other.
- Border leaf node and leaf node are not connected to each other.
- Spine nodes are not connected to each other.
- Super-spine nodes are not connected to each other.
- A leaf node marked as "multi-homed" must have an MCT neighbor.
- A leaf node marked as "single-homed" is not connected to other leaf nodes.
- A border leaf node marked as "multi-homed" must have an MCT neighbor.
- A border leaf node marked as "single-homed" is not connected to other border leaf nodes.
- Device role (such as leaf, border-leaf, spine, and super-spine) is validated for a given device platform type (for example, SLX 9840 cannot be added as a leaf).

Tip

The validation process reports any errors as a response to the **efa fabric device add** or **efa fabric configure** operations. You can use the **efa fabric error show** command to export these errors to a CSV file.

Example

This example creates the fabric.

```
efa fabric create --name --stage5
```

Example

This example adds a device to the fabric.

```
efa fabric device add--name stage5 --username admin --password password
--leaf 10.20.50.205,10.20.50.206,10.20.50.207 --spine 10.20.50.203,10.20.50.204
--three-stage-pod podA --super-spine
```

Example

This example adds multiple devices to the fabric.

```
efa fabric device add-bulk --name stage5 --username admin --password password
--leaf 10.20.50.205,10.20.50.206,10.20.50.207 --spine 10.20.50.203,10.20.50.204
--three-stage-pod podA --super-spine 10.20.50.201,10.20.50.202 --five-stage-pod podC
```

Example

This example configures the fabric topology.

```
efa fabric configure --name stage5
```

Provisioning Model to Migrate a 3-Stage Clos to 5-Stage Clos Fabric

Use the **`efa fabric migrate`** command.

```
efa fabric migrate
      --type {3-to-5-stage}
      --source-fabric <source-3-stage-clos-fabric>
      --destination-3-stage-leaf-spine-pod <3-stage-pod-name>
      --destination-3-stage-border-leaf-pod <border-leaf-pod-name>
      --super-spine-asn-block <super-spine-asn-range>
      --super-spine-peer-group <peer-group-name>


efa:root)root@admin01:~# efa fabric migrate -help
Migrate a 3-stage CLOS fabric to 5-stage CLOS fabric
Usage:
  efa fabric migrate [flags]
Flags:
   --type string                                    Type of migration [3-to-5-stage]
(default "3-to-5-stage")
   --source-fabric string                           Name of the 3-stage CLOS fabric to be
migrated
   --destination-3-stage-leaf-spine-pod string    Name of the 3-stage POD into which the
leaf and spine devices (of the 3-stage CLOS fabric)
                                                    need to be moved during migrate
   --destination-3-stage-border-leaf-pod string   Name of the 3-stage POD into which the
border-leaf devices (of the 3-stage CLOS fabric)
                                                    need to be moved during migrate
   --super-spine-asn-block string                 ASN block to be used by the super-spine
devices of the migrated 5-stage CLOS fabric
   --super-spine-peer-group string                Peer Group to be used by the spine
devices of the migrated 5-stage CLOS fabric
```

1. No or least traffic loss.
2. The command supports migration of a single fabric from 3-stage to 5-stage and not merging of multiple 3-stage Clos fabrics into a single 5-stage fabric.
3. If border-leaf is present / not-present in the fabric, then you must / must not provide the destination-border-leaf-pod.
4. If the border-leaf is connected to the spine devices in an existing 3-stage Clos fabric and you want the border-leaf to be connected to the spine even in the migrated 5-stage Clos fabric, the migration allows this.. Provide the same pod name for both the destination-3-stage-leaf-spine-pod and destination-3-stage-border-leaf-pod.
5. Migration is supported with multiple single-homed, multiple multi-homed border-leaf devices.

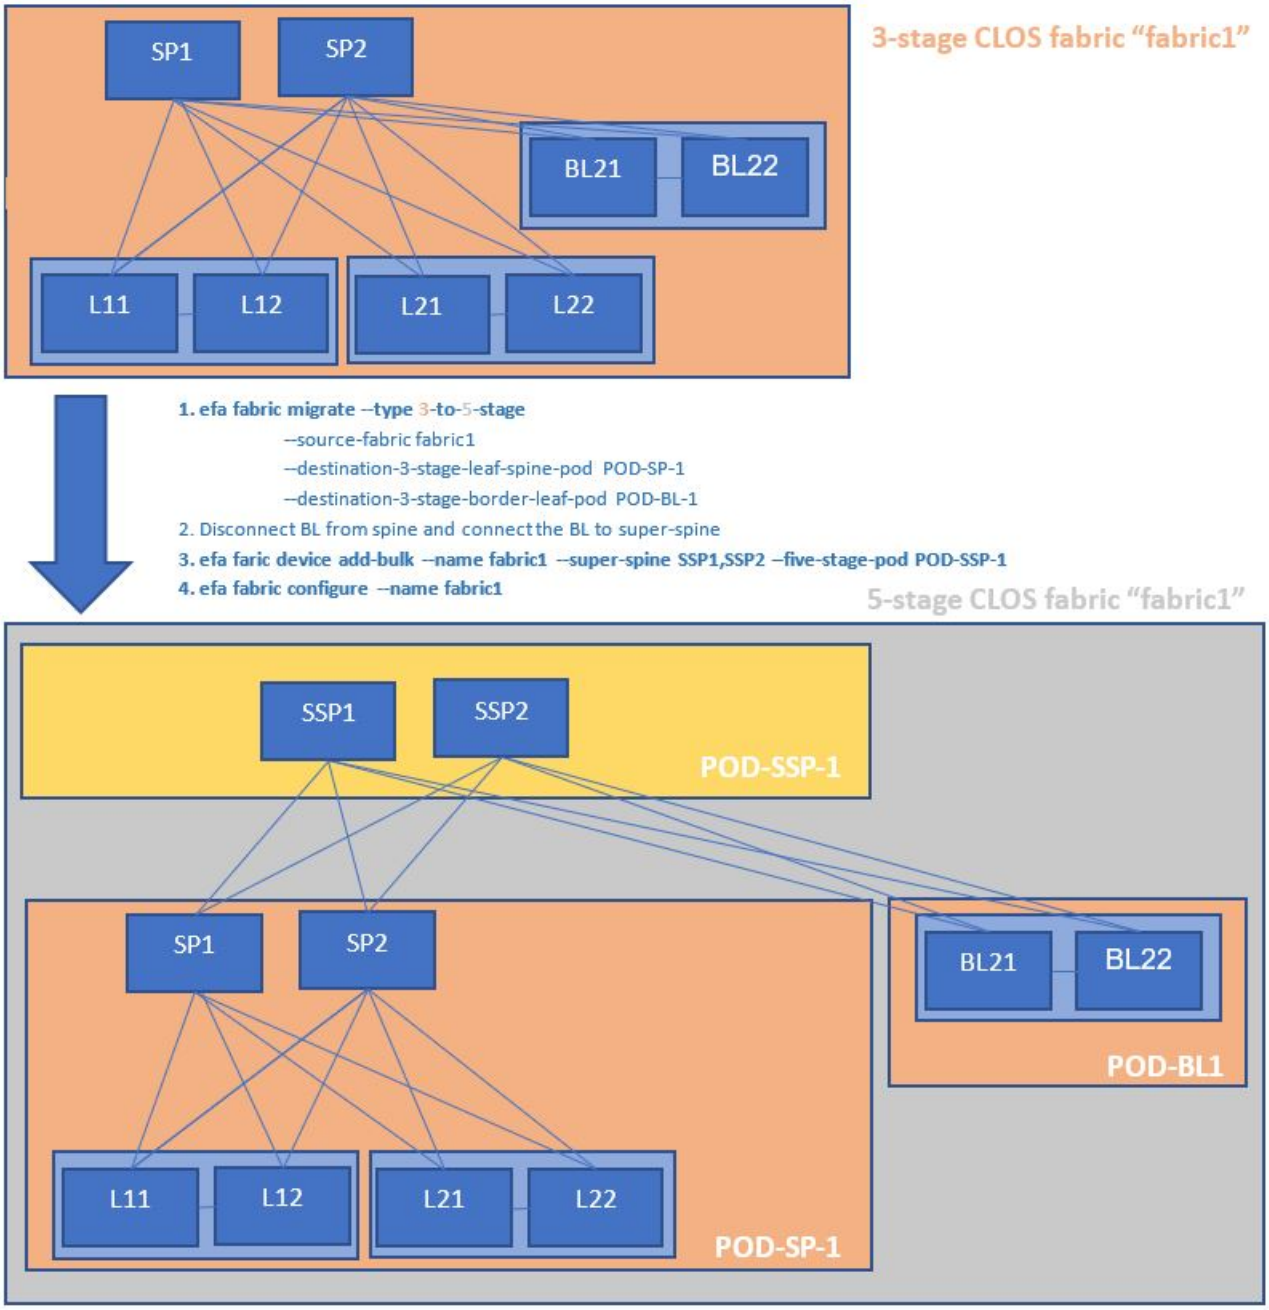

Migrate a 3-Stage Clos to 5-Stage Clos Fabric

You can migrate a 3-stage Clos to 5-stage Clos fabric.

About This Task

Complete the following tasks to migrate a 3-stage Clos to 5-stage Clos fabric

Procedure

1. Create a 3-Stage Clos Fabric on page 99
2. Migrate a 3-Stage Clos to 5-Stage Clos Fabric on page 99

3. Disconnect Border-leafs from Spine and Connect to Super-spine on page 101
4. Addition of Super-spine Devices to the Migrated 5-stage Clos Fabric on page 101
5. Configure the Migrated 5-stage Clos Fabric on page 101
6. Traffic Disruption during Fabric Configure on page 102
7. Verification of Fabric Underlay Configuration on the Migrated 5-stage Clos Fabric on page 103
8. Verification of Fabric Physical Underlay and Overlay Topology on the Migrated 5-stage Clos Fabric on page 106

Create a 3-Stage Clos Fabric

You can create a 3-stage Clos fabric.

About This Task

Use this procedure to create a 3-stage Clos fabric.

Procedure

Run the following command to create a 3-stage Clos fabric:

```
$ efa fabric create --name fabric1 --stage 3
$ efa fabric device add-bulk --name fabric1 --username admin --password password
                    --spine 10.17.112.221,10.17.112.222  --border-leaf 10.17.112.225-226
--leaf 10.17.112.223-224

$ efa fabric configure --name fabric1
$ efa fabric show --name fabric1
Fabric Name: fabric1, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success
+---------------+---+------+-------+-----------+------------+-----------+-----------
+--------+-----+----+
|  IP ADDRESS   |POD| HOST |  ASN  |   ROLE    |DEVICE STATE| APP
STATE | CONFIG GEN| PENDING| VTLB| LB |
|               |   | NAME |       |           |            |
| REASON    | CONFIGS| ID  | ID |
+---------------+---+------+-------+-----------+------------+-----------+-----------
+--------+-----+----+
| 10.17.112.221 |   | SP1  | 64512 |spine      |provisioned |cfg
in-sync| NA        | NA     | NA  | 1  |
| 10.17.112.222 |   | SP2  | 64512 |spine      |provisioned |cfg
in-sync| NA        | NA     | NA  | 1  |
| 10.17.112.223 |   | L11  | 65000 |leaf       |provisioned |cfg
in-sync| NA        | NA     | 2   | 1  |
| 10.17.112.224 |   | L12  | 65000 |leaf       |provisioned |cfg
in-sync| NA        | NA     | 2   | 1  |
| 10.17.112.225 |   | BL21 | 66000 |borderleaf |provisioned |cfg
in-sync| NA        | NA     | 2   | 1  |
| 10.17.112.226 |   | BL22 | 66000 |borderleaf |provisioned |cfg
in-sync| NA        | NA     | 2   | 1  |
+---------------+---+------+-------+-----------+------------+-----------+-----------
+--------+-----+----+
```

Migrate a 3-Stage Clos to 5-Stage Clos Fabric

When the 3-stage Clos fabric is migrated to a 5-stage Clos fabric,

1. Fabric stage is changed from “3” to “5”.
2. Fabric state is changed to “migrate-success” or “migrate-failed”.

3. Leaf and Spine devices will get the POD name = destination-3-stage-leaf-spine-pod.
4. Border Leaf devices will get the POD name = destination-3-stage-border-leaf-pod.
5. App State of all the leaf, border leaf and spine devices will be changed to “cfg-refreshed”.
6. Pending Configs of all the leaf, border leaf and spine devices will be changed to “BGP-C, BGP-D”.

Note
Brownfield configuration is not supported on the new devices added to the fabric in "migrate-success" state.

```
$ efa fabric migrate --type "3-to-5-stage" --source-fabric fabric1
             --destination-3-stage-leaf-spine-pod POD-SP-1 --destination-3-stage-border-
leaf-pod POD-BL-1

Fabric Name: fabric1, Fabric Description: , Fabric Stage: 5, Fabric Type: clos, Fabric
Status: migrate-success
+---------------+---------+------+------+----------+-------------+--------------
+-----------+------------------+------+----+
|  IP ADDRESS   |  POD    | HOST | ASN  |   ROLE
| DEVICE STATE|   APP STATE  | CONFIG GEN| PENDING CONFIGS  | VTLB | LB |
|               |         | NAME |      |          |
|              | REASON    |                 | ID   | ID |
+---------------+---------+------+------+----------+-------------+--------------
+-----------+------------------+------+----+
| 10.17.112.221 |POD-SP-1 | SP1  |64512 |spine     |
provisioned | cfg refreshed| NA        |SYSP-U,BGP-C,BGP-D| NA   | 1  |
| 10.17.112.222 |POD-SP-1 | SP2  |64512 |spine     |
provisioned | cfg refreshed| NA        |SYSP-U,BGP-C,BGP-D| NA   | 1  |
| 10.17.112.223 |POD-SP-1 | L11  |65000 |leaf      |
provisioned | cfg refreshed| NA        |SYSP-U,BGP-C,BGP-D| 2    | 1  |
| 10.17.112.224 |POD-SP-1 | L12  |65000 |leaf      |
provisioned | cfg refreshed| NA        |SYSP-U,BGP-C,BGP-D| 2    | 1  |
| 10.17.112.225 |POD-BL-1 | BL21 |66000 |borderleaf| provisioned
| cfg refreshed| NA        |SYSP-U,BGP-C,BGP-D| 2    | 1  |
| 10.17.112.226 |POD-BL-1 | BL22 |66000 |borderleaf| provisioned
| cfg refreshed| NA        |SYSP-U,BGP-C,BGP-D| 2    | 1  |
+---------------+---------+------+------+----------+-------------+--------------
+-----------+------------------+------+----+
```

The following diagram depicts a 5-stage topology after fabric migration:

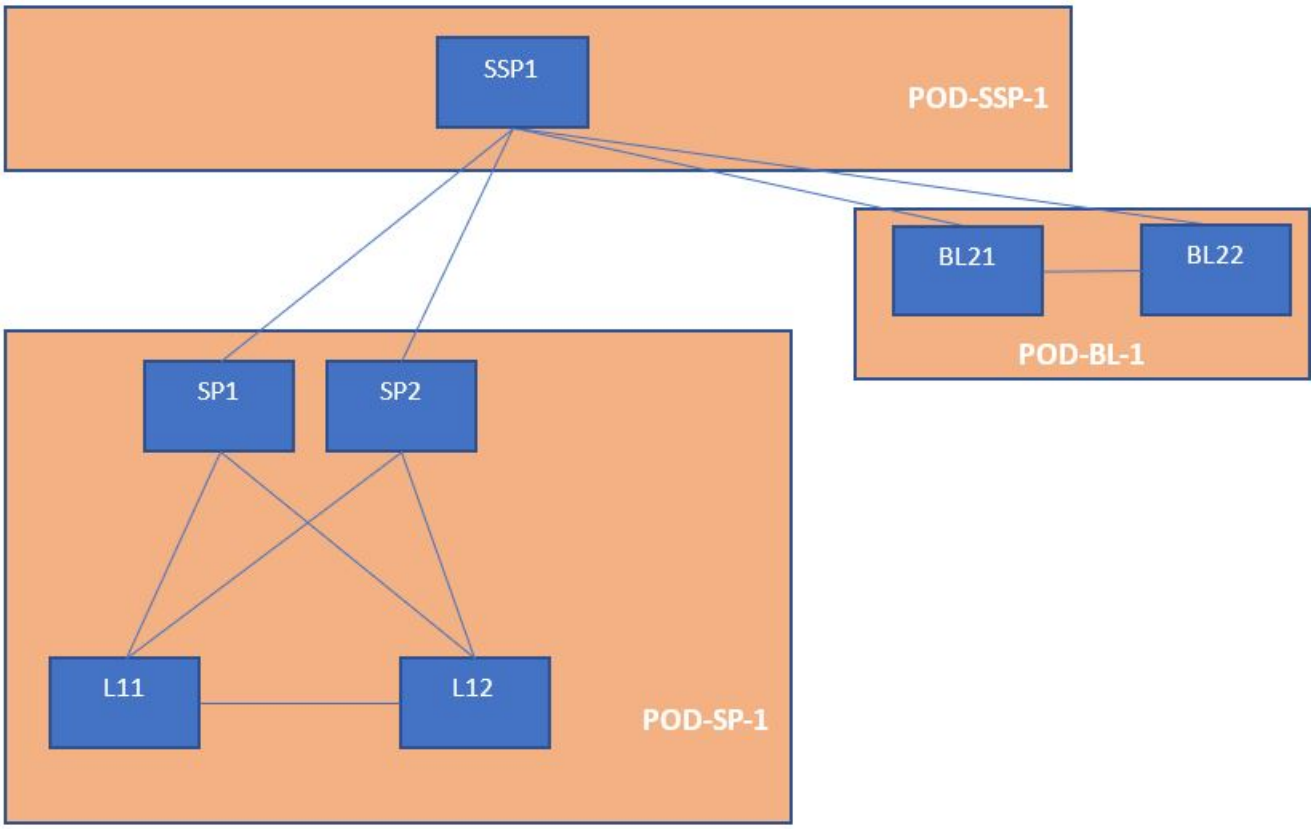

Disconnect Border-leafs from Spine and Connect to Super-spine

Disconnect the border-leaf devices, which were connected to spine devices in the 3-stage Clos fabric, from the spine devices, and then reconnect to the super-spine devices.

To keep the border-leaf devices connected to the spine devices (as done in the 3-stage Clos fabric) even in the 5-stage Clos fabric, do not disconnect the border-leaf devices from the spine devices, and then reconnect the border-leaf devices to the super-spine devices.

Addition of Super-spine Devices to the Migrated 5-stage Clos Fabric

Add the super-spine devices (which need to be part of the migrated fabric) to the migrated 5-stage Clos fabric.

```
$ efa fabric device add-bulk --name fabric1 --username admin --password password --super-spine 10.17.112.228   --five-stage-pod POD-SSP-1
```

Configure the Migrated 5-stage Clos Fabric

Configure the 5-stage Clos fabric with leaf, spine, super-spine, and border-leaf devices. The fabric topology validation is done during configuration of the fabric.

```
$ efa fabric configure --name fabric1
```

```
$ efa fabric show --name fabric1
Fabric Name: fabric1, Fabric Description: , Fabric Stage: 5, Fabric Type: clos, Fabric
Status: configure-success
+---------------+----------+-----+-------+-----------+-------------+-----------
+-----------+--------+-----+----+
|  IP ADDRESS   |   POD    |HOST |  ASN  |    ROLE
| DEVICE STATE| APP STATE | CONFIG GEN| PENDING| VTLB| LB |
|               |          |NAME |       |           |
|           | REASON    | CONFIGS| ID  | ID |
+---------------+----------+-----+-------+-----------+-------------+-----------
+-----------+--------+-----+----+
| 10.17.112.228 |POD-SSP-1 |SSP1 | 64769 | superspine|
provisioned |cfg in-sync| NA        | NA     | NA  | 1  |
| 10.17.112.225 |POD-BL-1  |BL21 | 66000 | borderleaf|
provisioned |cfg in-sync| NA        | NA     | 2   | 1  |
| 10.17.112.226 |POD-BL-1  |BL22 | 66000 | borderleaf|
provisioned |cfg in-sync| NA        | NA     | 2   | 1  |
| 10.17.112.221 |POD-SP-1  |SP1  | 64512 | spine
| provisioned |cfg in-sync| NA        | NA     | NA  | 1  |
| 10.17.112.222 |POD-SP-1  |SP2  | 64512 | spine
| provisioned |cfg in-sync| NA        | NA     | NA  | 1  |
| 10.17.112.223 |POD-SP-1  |L11  | 65000 | leaf
| provisioned |cfg in-sync| NA        | NA     | 2   | 1  |
| 10.17.112.224 |POD-SP-1  |L12  | 65000 | leaf
| provisioned |cfg in-sync| NA        | NA     | 2   | 1  |
+---------------+----------+-----+-------+-----------+-------------+-----------
+-----------+--------+-----+----+
```

Traffic Disruption during Fabric Configure

When you run the **`efa fabric configure`** command on a fabric in `migrate-success` state, the BGP sessions will be cleared on all the devices of the fabric in a phased manner, which is similar to the clearing of BGP sessions performed during the update of MD5 password on an active fabric followed by `efa fabric configure`.

Verification of Fabric Underlay Configuration on the Migrated 5-stage Clos Fabric

Complete the following configuration on SLX devices:

SPINE	SUPER SPINE
SP1# show running-config router bgp router bgp local-as 64512 capability as4-enable fast-external-fallover **neighbor POD-SP-1-leaf-group peer-group** neighbor POD-SP-1-leaf-group description To Leaf neighbor POD-SP-1-leaf-group bfd **neighbor POD-SSP-1-spine-group peer-group** neighbor POD-SSP-1-spine-group remote-as 64769 neighbor POD-SSP-1-spine-group description To SuperSpine neighbor POD-SSP-1-spine-group bfd neighbor 10.10.10.32 remote-as 65000 neighbor 10.10.10.32 peer-group POD-SP-1-leaf-group neighbor 10.10.10.34 remote-as 65000 neighbor 10.10.10.34 peer-group POD-SP-1-leaf-group neighbor 10.10.10.44 remote-as 65000 neighbor 10.10.10.44 peer-group POD-SP-1-leaf-group neighbor 10.10.10.46 remote-as 65000 neighbor 10.10.10.46 peer-group POD-SP-1-leaf-group neighbor 10.10.10.65 peer-group POD-SSP-1-spine-group address-family ipv4 unicast maximum-paths 8 graceful-restart ! address-family ipv6	SSP1# show running-config router bgp router bgp local-as 64769 capability as4-enable fast-external-fallover **neighbor POD-BL-1-leaf-group peer-group** neighbor POD-BL-1-leaf-group description To BorderLeaf neighbor POD-BL-1-leaf-group bfd **neighbor POD-SP-1-spine-group peer-group** neighbor POD-SP-1-spine-group remote-as 64512 neighbor POD-SP-1-spine-group description To Spine neighbor POD-SP-1-spine-group bfd neighbor 10.10.10.49 remote-as 66000 neighbor 10.10.10.49 peer-group POD-BL-1-leaf-group neighbor 10.10.10.51 remote-as 66000 neighbor 10.10.10.51 peer-group POD-BL-1-leaf-group neighbor 10.10.10.53 remote-as 66000 neighbor 10.10.10.53 peer-group POD-BL-1-leaf-group neighbor 10.10.10.55 remote-as 66000 neighbor 10.10.10.55 peer-group POD-BL-1-leaf-group neighbor 10.10.10.56 remote-as 66000 neighbor 10.10.10.56 peer-group POD-BL-1-leaf-group neighbor 10.10.10.58 remote-as 66000 neighbor 10.10.10.58 peer-group POD-BL-1-leaf-group neighbor 10.10.10.61 remote-as 66000 neighbor 10.10.10.61 peer-group POD-BL-1-leaf-group neighbor 10.10.10.63 remote-as 66000

```
unicast

 !

 address-family l2vpn
evpn

  graceful-
restart

  retain route-target
all

  neighbor POD-SSP-1-spine-group
encapsulation vxlan
  neighbor POD-SSP-1-spine-group
next-hop-unchanged
  neighbor POD-SSP-1-spine-group
enable-peer-as-check
  neighbor POD-SSP-1-spine-group
activate
  neighbor POD-SP-1-leaf-group
encapsulation vxlan
  neighbor POD-SP-1-leaf-group
next-hop-unchanged
  neighbor POD-SP-1-leaf-group
enable-peer-as-check
  neighbor POD-SP-1-leaf-group
activate
```

```
 neighbor 10.10.10.63 peer-group
POD-BL-1-leaf-group
 neighbor 10.10.10.64 peer-group
POD-SP-1-spine-group
 neighbor 10.10.10.67 peer-group
POD-SP-1-spine-group
 address-family ipv4
unicast

  maximum-paths
8

  graceful-
restart

 !

 address-family ipv6
unicast

 !

 address-family l2vpn
evpn

  graceful-
restart

  retain route-target
all

  neighbor POD-BL-1-leaf-group
encapsulation vxlan
  neighbor POD-BL-1-leaf-group
next-hop-unchanged
  neighbor POD-BL-1-leaf-group
enable-peer-as-check
  neighbor POD-BL-1-leaf-group
activate
  neighbor POD-SP-1-spine-group
encapsulation vxlan
  neighbor POD-SP-1-spine-group
next-hop-unchanged
  neighbor POD-SP-1-spine-group
enable-peer-as-check
  neighbor POD-SP-1-spine-group
activate
 !

!
```

LEAF

```
L11# show running-config router
bgp
router
bgp

 local-as
65000

 capability as4-
enable

 fast-external-
```

BORDER LEAF

```
BL21# show running-config router
bgp
router
bgp

 local-as
66000

 capability as4-
enable

 fast-external-
```

```
fallover

 neighbor POD-SP-1-spine-group
peer-group
 neighbor POD-SP-1-spine-group
remote-as 64512
 neighbor POD-SP-1-spine-group
description To Spine
 neighbor POD-SP-1-spine-group
bfd
 neighbor 10.10.10.41 peer-group
POD-SP-1-spine-group
 neighbor 10.10.10.43 peer-group
POD-SP-1-spine-group
 neighbor 10.10.10.45 peer-group
POD-SP-1-spine-group
 neighbor 10.10.10.47 peer-group
POD-SP-1-spine-group
 neighbor 10.20.20.6 remote-as
65000
 neighbor 10.20.20.6 next-hop-
self
 neighbor 10.20.20.6
bfd

 address-family ipv4
unicast

  network
172.31.254.73/32

  maximum-paths
8

  graceful-
restart

 !

 address-family ipv6
unicast

 !

 address-family l2vpn
evpn

  graceful-
restart

  neighbor POD-SP-1-spine-group
encapsulation vxlan
  neighbor POD-SP-1-spine-group
next-hop-unchanged
  neighbor POD-SP-1-spine-group
enable-peer-as-check
  neighbor POD-SP-1-spine-group
activate
 !

!
```

```
fallover

 neighbor POD-BL-1-spine-group
peer-group
 neighbor POD-BL-1-spine-group
description To Spine
 neighbor POD-BL-1-spine-group
bfd
 neighbor POD-SSP-1-spine-group
peer-group
 neighbor POD-SSP-1-spine-group
remote-as 64769
 neighbor POD-SSP-1-spine-group
description To SuperSpine
 neighbor POD-SSP-1-spine-group
bfd
 neighbor 10.10.10.48 peer-group
POD-SSP-1-spine-group
 neighbor 10.10.10.52 peer-group
POD-SSP-1-spine-group
 neighbor 10.10.10.60 peer-group
POD-SSP-1-spine-group
 neighbor 10.10.10.62 peer-group
POD-SSP-1-spine-group
 neighbor 10.20.20.8 remote-as
66000
 neighbor 10.20.20.8 next-hop-
self
 neighbor 10.20.20.8
bfd

 address-family ipv4
unicast

  network
172.31.254.106/32

  maximum-paths
8

  graceful-
restart

 !

 address-family ipv6
unicast

 !

 address-family l2vpn
evpn

  graceful-
restart

  neighbor POD-SSP-1-spine-group
encapsulation vxlan
  neighbor POD-SSP-1-spine-group
next-hop-unchanged
  neighbor POD-SSP-1-spine-group
enable-peer-as-check
  neighbor POD-SSP-1-spine-group
activate
```

	neighbor POD-BL-1-spine-group encapsulation vxlan neighbor POD-BL-1-spine-group next-hop-unchanged neighbor POD-BL-1-spine-group enable-peer-as-check neighbor POD-BL-1-spine-group activate ! !

Verification of Fabric Physical Underlay and Overlay Topology on the Migrated 5-stage Clos Fabric

After successful migration and configuration of fabric, run the following command to verify the physical, underlay, and overlay topology of the fabric:

```
efa fabric topology show {physical | underlay | overlay} --name fabric1
```

Operations Allowed on a Fabric in Migrate-failed State

The following operation is allowed on a fabric in `migrate-failed` state:

- Fabric Migrate

Operations Allowed on a Fabric in Migrate-success State

The following operations are allowed on a migrated fabric (fabric in `migrate-success` or `migrate-failed` state):

1. Fabric Clone
2. Fabric Delete
3. Fabric Device Add
4. Fabric Device Remove
5. Fabric Configure
6. Fabric Topology Show
7. Fabric Show

Operations not Allowed on a Fabric in Migrate-success and Migrate-failed State

- Run the **`efa fabric configure`** command to bring the fabric out of `migrate-success` state into `configure-success` state.
- Run the **`efa fabric migrate`** command to bring the fabric out of `migrate-failed` state into `migrate-success` state.
- Run the **`efa fabric configure`** command to bring the fabric out of `migrate-success` state into `configure-success` state.

The following operations are not allowed on a migrated fabric (fabric in `migrate-success` or `migrate-failed` state):

- Drift and Reconcile on a fabric device

- Fabric Setting update

Conditions Supporting Fabric Migration

The fabric migration is allowed in the following conditions:

1. Fabric is a 3-stage CLOS fabric
2. Fabric is in created state
3. Fabric is in configure-success state
4. Fabric is in migrate-failed state

Conditions Not Supporting Fabric Migration

The fabric migration is not allowed in the following conditions:

1. Fabric is of non-Clos type
2. Fabric is already a 5-stage Clos fabric
3. Fabric in configure-failed state
4. Fabric in migrate-success state
5. Fabric in settings-updated state
6. Fabric device dev-state is other than the cfg-in-sync or cfg-refreshed state

Supported Topology

1. Migration of a 3-stage Clos containing Leaf, Spine, and no Border Leaf to a 5-stage Clos.
2. Migration of a 3-stage Clos containing Leaf, Spine, and Border Leaf to a 5-stage Clos with Border Leaf continuing to be connected to Spine.
3. Migration of a 3-stage Clos containing Leaf, Spine, and Border Leaf to a 5-stage Clos with Border Leaf connected to Super-spine.

Non-Clos Small Data Center Overview

Support for Small DC Fabric offers CLI commands along with a REST API, similar to that of Clos Fabric.

Non-Clos fabric is supported on SLX 9150, SLX 9140, and SLX 9250 devices as follows:

- Single rack automation. Each rack consists of two node MCT pair.
- Multi-rack automation
- Multi-homed leaf (MCT)
- Overlay only automation
- Fabric topology view
- Fabric validation and troubleshooting

Supported Small Data Center Topologies

XCO supports small data center (non-Clos) fabrics.

XCO provides the following support for small data center fabrics on, SLX 9150, SLX 9250, and SLX 9740 devices:

- Single rack automation. Each rack consists of a two-node MCT pair.
- Multi-rack automation
- Multi-homed leaf (MCT)
- Overlay-only automation
- Fabric topology view
- Fabric validation and troubleshooting

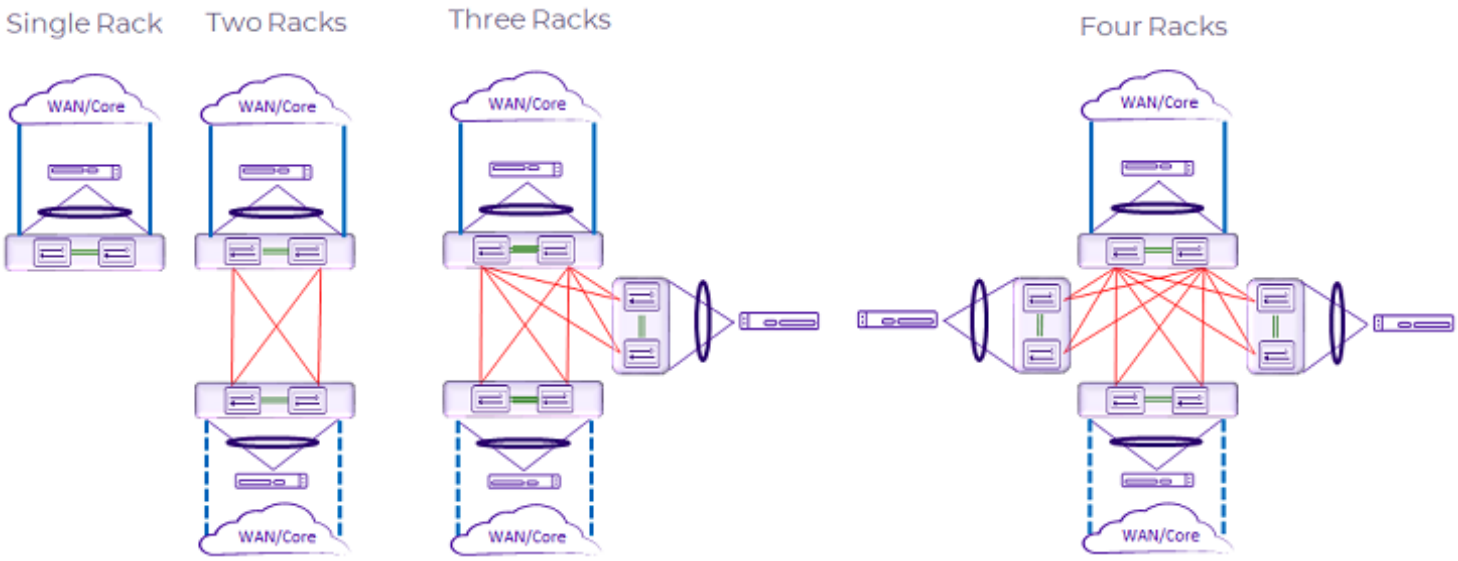

Figure 11: Supported small data center topologies

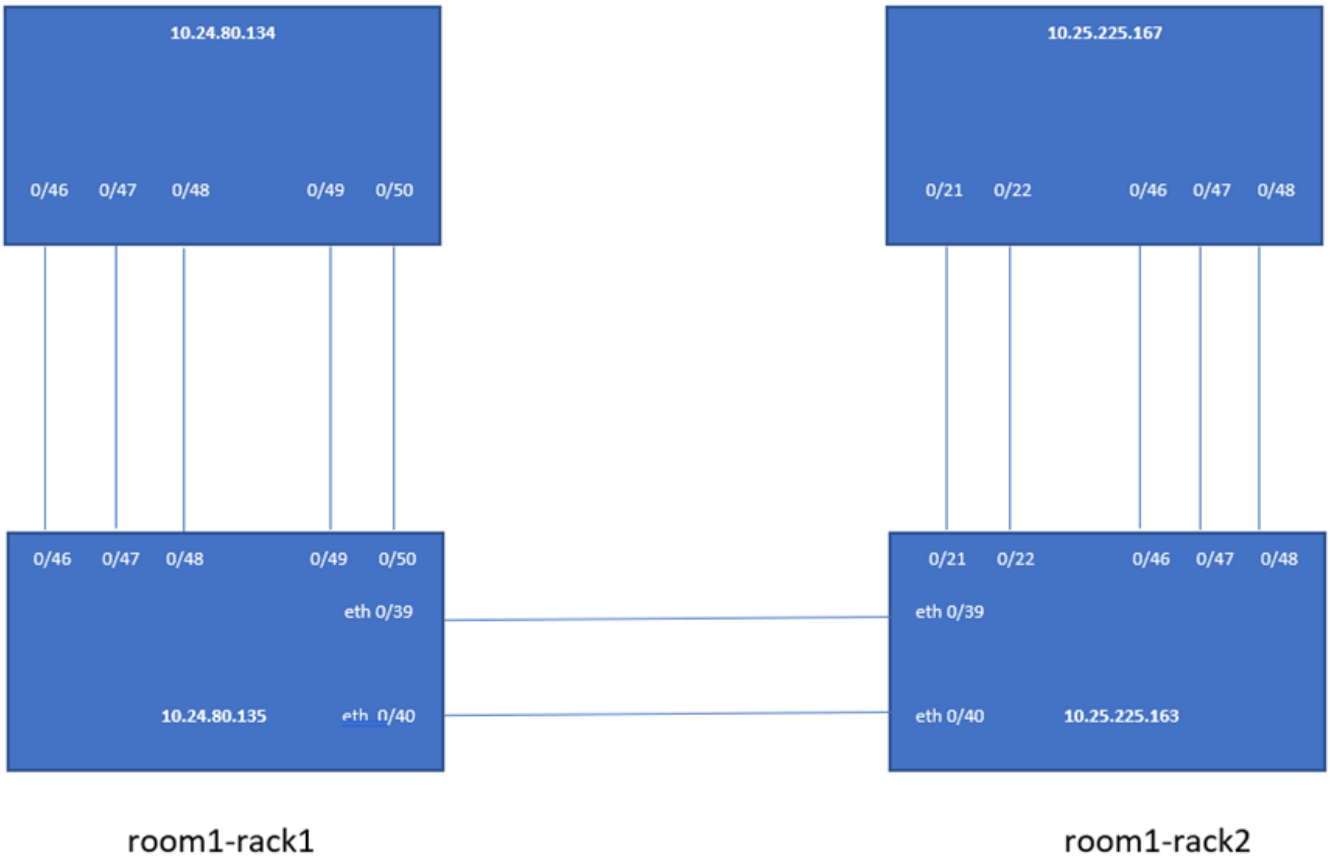

Figure 12: Multi-rack configuration example

Configure a Small Data Center Fabric

About This Task

Use this procedure to configure a small data center fabric.

Tip

If any devices in a fabric are in "admin-down" state, use of the following commands in that same fabric will not add or delete devices in the fabric: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

Procedure

1. Create a fabric.

```
efa fabric create
```

2. Add a device to the fabric.

```
efa fabric device add
```

   A device must be registered with Inventory Service before you can add it to a fabric. However, if you provide a user name and password when you run the command, then the devices are automatically registered with the Inventory Service. See the examples at the end of this procedure.

   You can add multiple devices by using the **efa fabric device add-bulk** command. If you choose to add multiple devices in bulk, ensure you perform the following operations first:

   - Run the **efa inventory device register --ip <list-of-device-ips>** command.
   - Run the **efa inventory device interface list --ip <device-ip>** command. In the output of the command, verify that the states of the port links are as you expected (in the Admin Status and Oper Status fields). If not, manually check the physical cabling and fix any issues. Then continue with the **efa fabric device add-bulk** operation.

3. Configure the fabric.

```
$ efa fabric configure
```

Tip
The validation process reports any errors as a response to the **efa fabric device add** or **efa fabric configure** operations. You can use the **efa fabric error show** command to export these errors to a CSV file.

Example

This example creates the fabric.

```
$ efa fabric create --name extr-fabric  --type non-clos
```

Example

This example adds a device to the fabric.

```
efa fabric device add --name extr-fabric --ip 10.x.x.x --rack room1-rack1
--username admin --password password
```

Example

This example adds multiple devices to the fabric.

```
$ efa fabric device add-bulk --name extr-fabric --rack room1-rack1
--ip 10.24.80.134,10.24.80.135 --rack room1-rack2 --ip 10.25.225.163,10.25.225.167
```

Example

This example configures the fabric.

```
efa fabric configure --name extr-fabric
```

Dynamic ICL in Small Data Center

Dynamic Inter-Chassis Link (ICL) in small data center (non-Clos) fabric dynamically identifies the ICL links for all the racks. In the latest version, you cannot specify the ICL ports manually.

Dynamic ICL in the non-Clos (small data center) fabric contains the following configuration changes:

1. Fabric settings CLI does not contain MCT ports and L3 backup port options.
2. There is no distinction between small data center and Clos MCT ports or LD-MCT ports in the backend. They do not have static configuration options and are identified by LLDP.
3. XCO running config does not have any references to the MCT ports.
4. Dynamic ICL configuration maintains backward compatibility for upgrade and downgrade operations for XCO.

Note

For small data center fabrics in EFA 2.5.5, MCT ICL ports are defined in the fabric settings; `rack-mct-ports` or `rack-ld-mct-ports`, and are the only ports used to form an ICL. Starting in EFA 2.6.0 and above, all the LLDP enabled ports that interconnect the MCT nodes are used for MCT ICL

Therefore, any ports, which were not previously defined in the fabric settings; `rack-mct-ports` or `rack-ld-mct-ports`, but are interconnecting the MCT nodes, will be automatically used for the MCT ICL after the upgrade from EFA 2.5.5 to EFA 2.6.0 and above.

Overview of Day-0 Operations for a Small Data Center Fabric

Day-0 operations consist of forming the fabric.

This table provides examples of the commands that you use to create a small data center (non-Clos) fabric with two SLX devices. For more information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Table 12: Day-0 operations

Operation	Command Example
Create a fabric	**`efa fabric create --name CNCF type non-clos`**
Enable backup routing	`efa fabric setting update --backup-routing-enable Yes --name CNCF`
Disable VLAN VNI auto-map	`efa fabric setting update --vni-auto-map No --name CNCF`
Add the first device	`efa fabric device add --ip 10.24.80.158 --hostname slx-a --rack pod1 --username admin --password password --name CNCF`

Table 12: Day-0 operations (continued)

Operation	Command Example
Add the second device	`efa fabric device add --ip 10.24.80.159 --hostname slx-b --rack pod1 --username admin --password password --name CNCF`
Configure the fabric	`efa fabric configure --name CNCF`

Configure Local Bias for Handling the LVTEP BUM Traffic

LVTEP is a Logical VTEP (Virtual Tunnel End Point) distributed across an MCT pair (Multi-Chassis Tunnel pair) with both devices configured with same VTEP IP. By default, the `Local Bias` is disabled for LVTEP BUM (Broadcast, Unicast, and Multicast) traffic handling.

About This Task

Follow this procedure to configure Local Bias.

- **When Local Bias is disabled**: The LVTEP BUM traffic handling is based on the DF (Designated Forwarder). One of the MCT devices becomes the DF for odd VLANs or BDs. The other MCT device becomes the DF for even VLANs or BDs. When VTEP on one of the MCT devices fails, the other MCT device must assigned itself as the DF for ALL VLANs or BDs.
- **When Local Bias is enabled**: Each MCT leaf device acts as a DF for all the VLANs or BDs. Each MCT leaf device can locally forward the BUM traffic towards the remote VTEP. This results in better traffic convergence when the LVTEP on one of the leaf devices goes down.

Procedure

To configure the Local Bias to handle the LVTEP BUM traffic, run the following command:

```
efa fabric setting update --name <fabric-name> --overlay-gateway-broadcast-local-bias {Yes|No}
```

- On a switching device, the `lvtep-broadcast-local-bias` configuration enables Local Bias to handle LVTEP BUM traffic.

  Note
  - The Local Bias configuration is not allowed if the cluster configuration already exists on the switching device.
  - You can provide the `overlay-gateway-broadcast-local-bias` fabric setting on a fabric before the fabric is configured.
  - You cannot modify the fabric settings on an existing deployments.

- The `lvtep-broadcast-local-bias` fabric settings apply on switches even though the fabric configured the feature with the supported SLX Platforms (SLX-9540, SLX-9640, SLX-9150, SLX-9250, SLX-9740, SLX-8520, SLX-8720, SLX-882, and SLX Firmware Version 20.4.3 and above).

If there is any drift in Global Configuration, the **`efa fabric show`** command displays that the "SYSP-C/U" configuration is pending.

- The `lvtep-broadcast-local-bias` configuration is applicable only for the dual-homed leaf or border-leaf devices and is not applicable for the single-homed leaf or border-leaf devices.
- The `lvtep-broadcast-local-bias` configuration is applicable for both Clos and non-Clos fabrics.
- The `lvtep-broadcast-local-bias` fabric setting can be enabled only when the `backup-routing-enable` fabric setting is enabled.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Example

Tho following example configures a Local Bias when you update fabric settings:

```
$ efa fabric setting update --name fs --overlay-gateway-broadcast-local-bias-enable yes
fs Fabric Update Successful



$ efa fabric setting show --advanced --name fs
+--------------------------------+-------------------------+
|              NAME              |          VALUE          |
+--------------------------------+-------------------------+
| Fabric Name                    | fs                      |
+--------------------------------+-------------------------+
| Link IP Range                  | 10.10.10.0/23           |
+--------------------------------+-------------------------+
| Loopback IP Range              | 172.31.254.0/24         |
+--------------------------------+-------------------------+
| Loopback Port Number           | 1                       |
+--------------------------------+-------------------------+
| VTEP Loopback Port Number      | 2                       |
+--------------------------------+-------------------------+
| Spine ASN Block                | 64512-64768             |
+--------------------------------+-------------------------+
| SuperSpine ASN Block           | 64769                   |
+--------------------------------+-------------------------+
| Leaf ASN Block                 | 65000-65534             |
+--------------------------------+-------------------------+
| Border Leaf ASN Block          | 66000-66100             |
+--------------------------------+-------------------------+
| P2P IP Type                    | numbered                |
+--------------------------------+-------------------------+
| Any cast MAC                   | 0201.0101.0101          |
+--------------------------------+-------------------------+
| IPV6 Any cast MAC              | 0201.0101.0102          |
+--------------------------------+-------------------------+
| MAC Aging Timeout              | 1800                    |
+--------------------------------+-------------------------+
| MAC Aging Conversational       | 300                     |
| Timeout                        |                         |
+--------------------------------+-------------------------+
| MAC Move Limit                 | 20                      |
+--------------------------------+-------------------------+
| Duplicate MAC Timer            | 5                       |
+--------------------------------+-------------------------+
| Duplicate MAC Timer MAX Count  | 3                       |
+--------------------------------+-------------------------+
```

```
| BFD Enable                       | Yes                      |
+----------------------------------+--------------------------+
| BFD Tx                           | 300                      |
+----------------------------------+--------------------------+
| BFD Rx                           | 300                      |
+----------------------------------+--------------------------+
| BFD Multiplier                   | 3                        |
+----------------------------------+--------------------------+
| MaxPaths                         | 8                        |
+----------------------------------+--------------------------+
| AllowAsIn                        | 0                        |
+----------------------------------+--------------------------+
| MTU                              | 9216                     |
+----------------------------------+--------------------------+
| IPMTU                            | 9100                     |
+----------------------------------+--------------------------+
| MCT Link IP Range                | 10.20.20.0/24            |
+----------------------------------+--------------------------+
| MCT PortChannel                  | 64                       |
+----------------------------------+--------------------------+
| LACP Timeout                     | long                     |
+----------------------------------+--------------------------+
| Control Vlan                     | 4090                     |
+----------------------------------+--------------------------+
| Control VE                       | 4090                     |
+----------------------------------+--------------------------+
| Leaf PeerGroup                   | spine-group              |
+----------------------------------+--------------------------+
| Spine PeerGroup                  | leaf-group               |
+----------------------------------+--------------------------+
| SuperSpine PeerGroup             | spine-group              |
+----------------------------------+--------------------------+
| Configure Overlay Gateway        | Yes                      |
+----------------------------------+--------------------------+
| VNI Auto Map                     | Yes                      |
+----------------------------------+--------------------------+
| Backup Routing Enable            | No                       |
+----------------------------------+--------------------------+
| Backup Routing IPv4 Range        | 10.40.40.0/24            |
+----------------------------------+--------------------------+
| Backup Routing IPv6 Range        | fd40:4040:4040:1::/120   |
+----------------------------------+--------------------------+
| Optimized Replication Enable     | No                       |
+----------------------------------+--------------------------+
| MDT Group IPv4 Range             | 239.0.0.0/8              |
+----------------------------------+--------------------------+
| Default MDT Group IPv4 address   | 239.1.1.1                |
+----------------------------------+--------------------------+
| MD5 Password Enable              | No                       |
+----------------------------------+--------------------------+
| MD5 Password                     |                          |
+----------------------------------+--------------------------+
| BGP Dynamic Peer Listen Limit    | 100                      |
+----------------------------------+--------------------------+
| Overlay Gateway Broadcast        | Yes                      |
| Local Bias Enable                |                          |
+----------------------------------+--------------------------+
```

The following is an example configuration on switch devices:

```
Rack1-Device1# show running-config
overlay-gateway
overlay-gateway fabric1
 ip interface Loopback 2
 map vni auto
 activate
!

Rack1-Device1# show running-config
interface Loopback  2
interface Loopback 2
 ip address 172.31.254.34/32
 no shutdown
!

Rack1-Device1# show running-config
lvtep-broadcast-local-bias
lvtep broadcast-local-bias
```

```
Rack1-Device2# show running-config
overlay-gateway
overlay-gateway fabric1
 ip interface Loopback 2
 map vni auto
 activate
!

Rack1-Device2# show running-config
interface Loopback  2
interface Loopback 2
 ip address 172.31.254.34/32
 no shutdown
!

Rack1-Device2# show running-config
lvtep-broadcast-local-bias
lvtep broadcast-local-bias
```

IP Multicast Fabric Provisioning

IP Multicast Fabric Overview

When multicast traffic is sent over unicast tunnels, ingress replication is done for each remote VTEP node. IP multicast fabric enables IP fabric to distribute BUM (Broadcast, Unknown Unicast, and Multicast Overlay) traffic using multicast VxLAN tunnels established over underlay fabric links.

Multicast Vxlan tunnels use Protocol Independent Multicast - Source Specific Multicast (PIM-SSM) and Multicast Distribution Tree (MDT) to deliver traffic effectively while minimizing packet replication in the fabric.

When multicast fabric is configured, a default MDT is created using PIM-SSM protocol running on fabric links and all the EVPN domain (VLANs and BDs) traffic is routed using the default tree.

The following figures show Clos topology for VxLAN unicast and multicast tunnels.

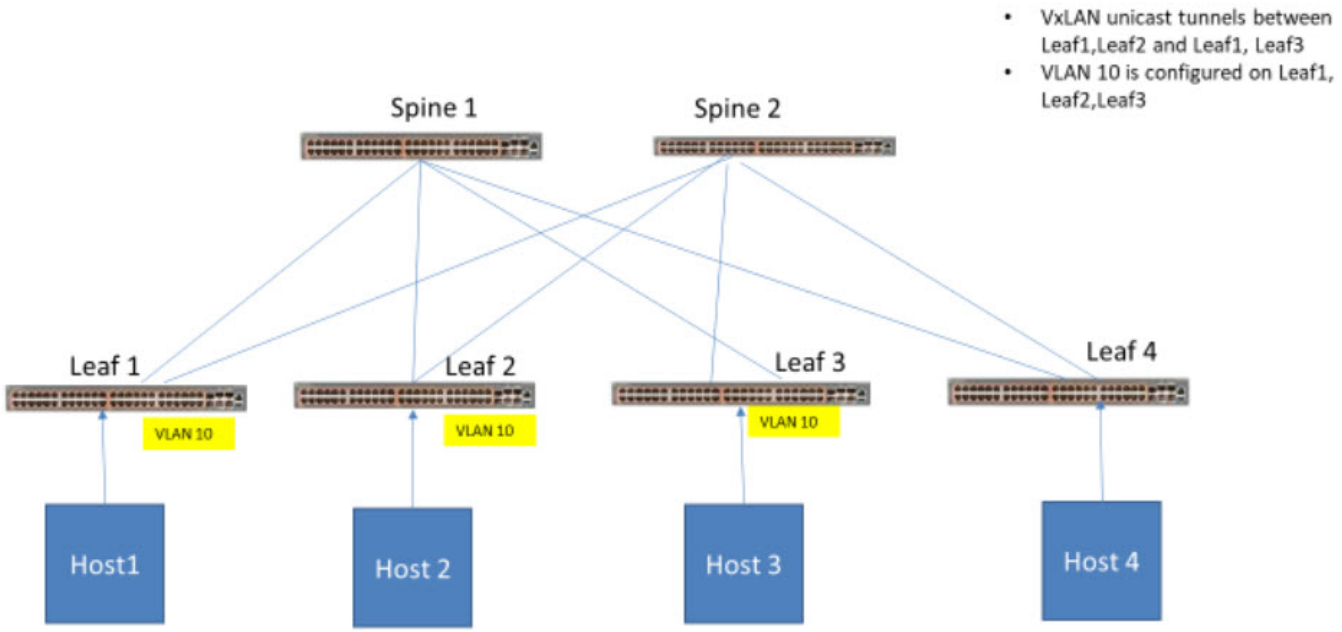

Figure 13: Clos topology with VxLAN unicast tunnels

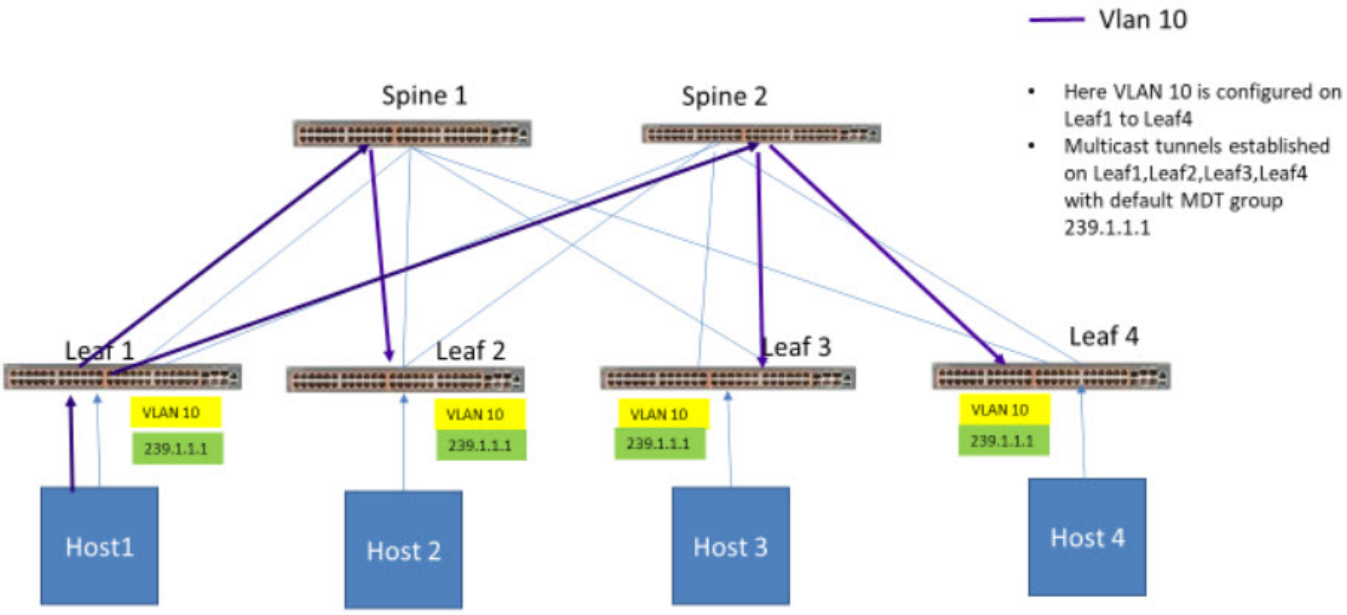

Figure 14: Clos topology with multicast tunnels

Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) protocol detects faults between two forwarding engines.

When fabric is created, BFD is enabled by default along with fabric links and BGP neighbors. The following example shows BFD configuration settings.

```
# efa fabric setting show --name clos_fabric --advanced

+--------------------------------+------------------------+
|              NAME              |         VALUE          |
+--------------------------------+------------------------+
| Fabric Name                    | default                |
+--------------------------------+------------------------+
| Link IP Range                  | 10.10.10.0/23          |
+--------------------------------+------------------------+
| Loopback IP Range              | 172.31.254.0/24        |
+--------------------------------+------------------------+
| Loopback Port Number           | 1                      |
+--------------------------------+------------------------+
| VTEP Loopback Port Number      | 2                      |
+--------------------------------+------------------------+
| Spine ASN Block                | 64512-64768            |
+--------------------------------+------------------------+
| SuperSpine ASN Block           | 64769                  |
+--------------------------------+------------------------+
| Leaf ASN Block                 | 65000-65534            |
+--------------------------------+------------------------+
| Border Leaf ASN Block          | 66000-66100            |
+--------------------------------+------------------------+
| P2P IP Type                    | numbered               |
+--------------------------------+------------------------+
| Any cast MAC                   | 0201.0101.0101         |
+--------------------------------+------------------------+
| IPV6 Any cast MAC              | 0201.0101.0102         |
+--------------------------------+------------------------+
| MAC Aging Timeout              | 1800                   |
+--------------------------------+------------------------+
| MAC Aging Conversational       | 300                    |
| Timeout                        |                        |
+--------------------------------+------------------------+
| MAC Move Limit                 | 20                     |
+--------------------------------+------------------------+
| Duplicate MAC Timer            | 5                      |
+--------------------------------+------------------------+
| Duplicate MAC Timer MAX Count  | 3                      |
+--------------------------------+------------------------+
| BFD Enable                     | Yes                    |
+--------------------------------+------------------------+
| BFD Tx                         | 300                    |
+--------------------------------+------------------------+
| BFD Rx                         | 300                    |
+--------------------------------+------------------------+
| BFD Multiplier                 | 3                      |
+--------------------------------+------------------------+
| BGP MultiHop                   | 2                      |
+--------------------------------+------------------------+
| MaxPaths                       | 8                      |
+--------------------------------+------------------------+
| AllowAsIn                      | 0                      |
+--------------------------------+------------------------+
| MTU                            | 9216                   |
+--------------------------------+------------------------+
```

```
| IPMTU                          | 9100                    |
+--------------------------------+-------------------------+
| MCT Link IP Range              | 10.20.20.0/24           |
+--------------------------------+-------------------------+
| MCT PortChannel                | 64                      |
+--------------------------------+-------------------------+
| LACP Timeout                   | long                    |
+--------------------------------+-------------------------+
| Control Vlan                   | 4090                    |
+--------------------------------+-------------------------+
| Control VE                     | 4090                    |

…Skipped
```

Fabric Settings to Update BGP MD5 Password, BGP Dynamic Peer Listen Limit, and Single Rack Deployment

Use the fabric settings to update BGP MD5 password using `--md5-password` and `--md5-password-enable`. Update the BGP dynamic peer listen limit using `--bgp-dynamic-peer-listen-limit`. A setting to denote a single rack or multi-rack deployment `--single-rack-deployment` is added under fabric setting. You can update these settings after fabric is configured.

For details on BGP MD5 Authentication, see *ExtremeCloud Orchestrator Security Guide, 3.2.1*.

The following is an example of output from the **`Rack Low Density L3 backup`** and **`efa fabric setting show --name --advanced`** command.

```
| Rack Low Density L3 backup     | 0/30                          |
| port (not applicable to        |                               |
| SLX-9250)                      |                               |
+--------------------------------+-------------------------------+
| Rack Low Density MCT Ports     | 0/19,0/22                     |
+--------------------------------+-------------------------------+
| Configure Overlay Gateway      | Yes                           |
+--------------------------------+-------------------------------+
| VNI Auto Map                   | Yes                           |
+--------------------------------+-------------------------------+
| Backup Routing Enable          | Yes                           |
+--------------------------------+-------------------------------+
| Backup Routing IPv4 Range      | 10.40.40.0/24                 |
+--------------------------------+-------------------------------+
| Backup Routing IPv6 Range      | fd40:4040:4040:1::/120        |
+--------------------------------+-------------------------------+
| MD5 Password Enable            | Yes                           |
+--------------------------------+-------------------------------+
| MD5 Password                   | $9$9Yc+EqhoseRMPh9fDhk9TQ==   |
+--------------------------------+-------------------------------+
| BGP Dynamic Peer Listen Limit  | 101                           |
+--------------------------------+-------------------------------+
```

```
# efa fabric setting show  --name fabric1 --advanced
+--------------------------------+-------------------------------+
|              NAME              |             VALUE             |
+--------------------------------+-------------------------------+
| Fabric Name                    | fabric1                       |
```

```
+--------------------------------+-----------------------------+
| Link IP Range                  | 10.10.10.0/23               |
+--------------------------------+-----------------------------+
| Loopback IP Range              | 172.31.254.0/24             |
+--------------------------------+-----------------------------+
| Loopback Port Number           | 1                           |
+--------------------------------+-----------------------------+
| VTEP Loopback Port Number      | 2                           |
+--------------------------------+-----------------------------+
| Rack ASN Block                 | 4200000000-4200065534       |
+--------------------------------+-----------------------------+
| Rack Border Leaf ASN Block     | 4200065535-4200065635       |
+--------------------------------+-----------------------------+
| Single Rack Deployment         | Yes                         |
+--------------------------------+-----------------------------+
| P2P IP Type                    | numbered                    |
+--------------------------------+-----------------------------+
| Any cast MAC                   | 0201.0101.0101              |
+--------------------------------+-----------------------------+
| IPV6 Any cast MAC              | 0201.0101.0102              |
+--------------------------------+-----------------------------+
| MAC Aging Timeout              | 1800                        |
+--------------------------------+-----------------------------+
| MAC Aging Conversational       | 300                         |
| Timeout                        |                             |
+--------------------------------+-----------------------------+
| MAC Move Limit                 | 20                          |
+--------------------------------+-----------------------------+
| Duplicate MAC Timer            | 5                           |
+--------------------------------+-----------------------------+
| Duplicate MAC Timer MAX Count  | 3                           |
+--------------------------------+-----------------------------+
| BFD Enable                     | Yes                         |
+--------------------------------+-----------------------------+
| BFD Tx                         | 400                         |
+--------------------------------+-----------------------------+
| BFD Rx                         | 400                         |
+--------------------------------+-----------------------------+
| BFD Multiplier                 | 5                           |
+--------------------------------+-----------------------------+
| BGP MultiHop                   | 4                           |
+--------------------------------+-----------------------------+
| MaxPaths                       | 8                           |
+--------------------------------+-----------------------------+
| AllowAsIn                      | 0                           |
+--------------------------------+-----------------------------+
| MTU                            | 9216                        |
+--------------------------------+-----------------------------+
| IPMTU                          | 9100                        |
+--------------------------------+-----------------------------+
| MCT Link IP Range              | 10.20.20.0/24               |
+--------------------------------+-----------------------------+
| MCT PortChannel                | 64                          |
+--------------------------------+-----------------------------+
| LACP Timeout                   | short                       |
+--------------------------------+-----------------------------+
| Control Vlan                   | 4090                        |
+--------------------------------+-----------------------------+
| Control VE                     | 4090                        |
+--------------------------------+-----------------------------+
| Rack L3 Backup IP Range        | 10.30.30.0/24               |
+--------------------------------+-----------------------------+
| Rack Underlay EBGP Peer Group  | underlay-ebgp-group         |
+--------------------------------+-----------------------------+
```

```
| Rack Overlay EBGP Peer Group   | overlay-ebgp-group          |
+--------------------------------+-----------------------------+
| Rack L3 backup port            | 0/48                        |
+--------------------------------+-----------------------------+
| Rack MCT Ports                 | 0/46,0/47                   |
+--------------------------------+-----------------------------+
```

Configure an IP Multicast Fabric

You can configure an IP multicast fabric.

About This Task

Use this procedure to configure an IP multicast fabric.

Tip

If any devices in a fabric are in "admin-down" state, use of the following commands in that same fabric will not add or delete devices in the fabric: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

Procedure

1. Create a Clos fabric.

```
# efa fabric create --name clos_fabric --type clos --stage 3
```

Note

Optimized replication is not supported on small data center fabric.

2. Enable multicast fabric settings.

```
# efa fabric setting update --optimized-replication-enable yes --name clos_fabric
```

3. (Optional) Override the default MDT group and group range.

```
# efa fabric setting update --name clos_fabric --mdtgroup-range <A.B.C.D/L> --default-
mdtgroup <A.B.C.D>

# efa fabric setting update --name clos_fabric --mdtgroup-range 239.0.0.0/8 --default-
mdtgroup 239.1.1.1
```

4. Verify the fabric settings.

```
# efa fabric setting show --name clos_fabric --advanced

+--------------------------------+-------------------------+
| Optimized Replication Enable   | Yes                     |
+--------------------------------+-------------------------+
| MDT Group IPv4 Range           | 238.0.0.0/8             |
+--------------------------------+-------------------------+
| Default MDT Group IPv4 address | 238.1.1.1               |
+--------------------------------+-------------------------+
```

5. Add devices to the fabric.

```
# efa fabric device add-bulk --name clos_fabric --leaf Leaf1IP,Leaf2IP,Leaf3IP,Leaf4IP
--spine Spine1IP,Spine2IP --username admin --password password
```

6. Configure the fabric.

```
# efa fabric configure -name clos_fabric
```

7. Verify the fabric configuration.

```
# efa fabric show-config --name clos_fabric --advanced
```

   - All underlay configuration and overlay configurations are pushed to the devices and underlay topology is operational.
   - All BGP connections between leaf and spine nodes are established and neighbors are reachable.
   - Basic overlay configuration with optimized replication is configured.
   - All device configurations are applied to the devices in fabric. For more information, see Device Configuration on page 121.

8. Create a tenant and endpoint group to bring up the multicast tunnels with leaf nodes.

```
# efa tenant create --name tenant1 --l2-vni-range 10002-14190 --
l3-vni-range 14191-14200 --vrf-count 10 --vlan-range 2-4090 --port
Leaf1IP[0/12-16],Leaf2IP[0/12-16], Leaf3IP[0/12-16],Leaf4IP[0/12-16] --description
Subscriber1

# efa tenant epg create --name epg1 --tenant tenant1 --port
Leaf1IP[0/15],Leaf2IP[0/15],Leaf3IP[0/16 --switchport-mode trunk --ctag-range 100
```

Device Configuration

When IP multicast fabric is enabled, the following device configurations are pushed to all devices in the fabric.

- `router pim` for default VRF is enabled on all nodes.
- `ip prefix-list` is configured on all nodes with the `mdt-range` specified in fabric settings or default range.
- Under `router-pim` mode, PIM-SSM is enabled for all nodes with range specified in `ip prefix-list`.
- Under `interface` mode, `PIM sparse` mode is enabled for all fabric links.
- Under overlay-gateway, `optimized replication` is enabled on all leaf nodes.
- Under `optimized replication` mode, `underlay-default-mdtgroup` is configured to default value specified in fabric settings on all leaf nodes.

Configure Drift and Reconcile on Multicast Fabric

You can configure drift and reconcile on multicast fabric.

About This Task

Use this procedure to configure drift and reconcile.

Procedure

1. To configure drift and reconcile on multicast fabric, run the following command:
```
# efa fabric debug device drift --ip A.B.C.D --name dni --reconcile
```

Note

Any drift in Router PIM, IP prefix list, and overlay-gateway XCO configuration compared to the configured device is detected and reconciled.

2. To configure drift and reconcile of all services on a device, run the following command:
```
# efa inventory drift-reconcile execute --ip A.B.C.D --reconcile
```

View Fabric Details

You can use several commands to view the details of topologies and configuration in your fabric.

Table 13: Fabric show commands

Command	Description
`efa fabric topology show overlay`	Shows the overlay (VxLAN tunnels) connectivity of devices in a fabric.
`efa fabric topology show physical`	Shows the physical topology (LLDP neighbors) connectivity of devices in a fabric.
`efa fabric topology show underlay`	Shows the underlay (BGP neighbors) connectivity of devices in a fabric.
`efa fabric show`	Displays fabric details.
`efa fabric show-config`	Displays fabric configuration details for the specified role (leaf, spine, super-spine, border leaf) or IP address.
`efa fabric show summary`	Displays a summary of all fabrics or of the specified fabric.

The following is an example of output from the **`efa fabric topology show overlay`** command.

```
efa fabric topology show overlay --name fabric1
+------------+--------------------------+--------------------------+----------------
+---------------+-------------+------------+----------
| OVERLAY    |      SOURCE LEAF IP      |    DESTINATION LEAF IP    |
SOURCE VTEP IP | DESTINATION   | OVERLAY     | OVERLAY    | OVERLAY
  ECAP TYPE
VTEP IP         ADMIN STATE   OPER STATE   BFD STATE
+------------+--------------------------+--------------------------+----------------
+---------------+-------------+------------+----------
| vxlan      | 10.25.225.11,10.25.225.46 | 10.24.85.76,10.24.85.74   | 172.31.254.86
| 172.31.254.81 | up          | up         | down
```

```
| vxlan      | 10.25.225.11,10.25.225.46 | 10.24.80.134,10.24.80.135 | 172.31.254.86
| 172.31.254.83 | up          | up          | down
| vxlan      | 10.24.85.76,10.24.85.74   | 10.25.225.11,10.25.225.46 | 172.31.254.81
| 172.31.254.86 | up          | up          | down
| vxlan      | 10.24.85.76,10.24.85.74   | 10.24.80.134,10.24.80.135 | 172.31.254.81
| 172.31.254.83 | up          | up          | down
| vxlan      | 10.24.80.134,10.24.80.135 | 10.25.225.11,10.25.225.46 | 172.31.254.83
| 172.31.254.86 | up          | up          | down
| vxlan      | 10.24.80.134,10.24.80.135 | 10.24.85.76,10.24.85.74   | 172.31.254.83
| 172.31.254.81 | up          | up          | down
+------------+---------------------------+---------------------------+----------------
+---------------+-------------+------------+----------
```

The following is an example of output from the **efa fabric show** command.

```
efa fabric topology show overlay --name fabric1
+--------------+-------+-------------+-----------+------+--------------+-------------
+------------+----------+------+-------
| IP ADDRESS   | RACK  |  HOST NAME  | ASN       | ROLE | DEVICE
STATE | APP STATE   | CONFIG     | PENDING  | VTLB | LB ID

GEN REASON   CONFIGS    ID
+--------------+-------+-------------+-----------+------+--------------+-------------
+------------+----------+------+-------
| 10.24.51.131 | rack2 | Freedom-07  | 4200000001| leaf | provisioned
| cfg in-sync | NA         | NA       | 2    | 1
| 10.25.255.58 | rack2 | Freedom-04  | 4200000001| leaf | provisioned
| cfg in-sync | NA         | NA       | 2    | 1
| 10.24.51.135 | rack1 | Freedom-06  | 4200000000| leaf | provisioned
| cfg in-sync | NA         | NA       | 2    | 1
| 10.24.48.131 | rack1 | Freedom-05  | 4200000002| leaf | provisioned
| cfg in-sync | NA         | NA       | 2    | 1
```

Edit Fabric Settings

You can edit certain fabric settings after the fabric configuration.

Changes made in the fabric settings are displayed in the **efa fabric show** command output. The **efa fabric show** command output marks the changed settings as `Updated Fabric Settings` and corresponding modification codes are displayed.

```
efa fabric show output:
...
Updated Fabric Settings: BGP-LL
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password , BFD-RX - Bfd Rx
Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd Timer,BFD-ENABLE - Enable Bfd, BGP-
MULTIHOP - BGP ebgp multihop, P2PLR - Point-to-Point Link Range, MCTLR - MCT Link Range,
LOIP - Loopback IP Range..
```

- Changes are not reflected in `app-state`, `drc-drift`, and `drc-reconcile show` outputs.
- Changes are pushed into SLX only after you run the **efa fabric configure** command.
- After successful execution of the **efa fabric configure** command, the **efa fabric show** command clears the fabric status and fabric settings.

Edit Fabric IP Range Settings

- You can edit the following active fabric IP range settings even after the devices are added in fabric:
  - `--loopback-ip-range`
  - `--p2p-link-range`
  - `--mctlink-ip-range`
- You can generate increased IP addresses or IP-Pairs and update them in the available IP-Pool. There are no updates to the used IP-pair. For example,
  - Old IP range: 10.10.10.1/24
  - New IP range: 10.10.10.1/23
  - Increased IP range: 10.10.10.1/23 - 10.10.10.1/24
- The increased IP addresses are available for use when new devices are added into the fabric.
- You can edit l`loopback-ip-range`, `p2p-link-range`, and `mctlink-ip-range` settings. Within the same network, ensure that the prefix mask length is lower than the configured prefix mask length.

```
'efa fabric show' output sample:
...
Fabric Name: fs, Fabric Description: , Fabric Type: non-clos, Fabric Status: settings-
updated
Updated Fabric Settings: MCTLR,P2PLR,LOIP
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password , BFD-RX - Bfd Rx
Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd Timer,BFD-ENABLE - Enable Bfd,
BGP-MULTIHOP - BGP ebgp multihop, P2PLR - Point-to-Point Link Range, MCTLR - MCT Link
Range, LOIP - Loopback IP Range
```

Edit Fabric BFD Settings

- You can edit the following active fabric BFD settings even after the devices are added in fabric:
  - `--bfd-enable <yes|no>`
  - `--bfd-tx <val>`
  - `---bfd-rx <val>`
  - `--bfd-multiplier<val>`

  The following is an example output of the **`efa fabric show`** command:

```
Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: settings-updated

Updated Fabric Settings: BFD-ENABLE,BFD-TX,BFD-RX,BFD-MULTIPLIER

FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password, BFD-RX - Bfd Rx
Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd Timer,BFD-ENABLE - Enable Bfd,
BGP-MULTIHOP - BGP ebgp multihop, P2PLR - Point-to-Point Link Range, MCTLR - MCT Link
Range, LOIP - Loopback IP Range
```

- The **`efa fabric configure`** command generates BFD recipe for router BGP, interface links, and MCT port-channel interfaces, and then pushes the configuration to the SLX devices.

Fabric Settings in Active Fabric: Small Data Center Fabric

Cannot be Updated	Can be Updated	Invalid Settings for Small Data Center Fabric

`--rack-l3-backup-ip-range`
`--loopback-port-number`
`--vtep-loopback-port-number`
`--rack-asn-block`
`--rack-border-leaf-asn-block`
`--anycast-mac-address`
`--ipv6-anycast-mac-address`
`--mac-aging-timeout`
`--mac-aging-conversation-timeout`
`--mac-move-limit`
`--duplicate-mac-timer`
`--duplicate-mac-timer-max-count`
`--max-paths`
`--allow-as-in`
`--mtu`
`--ip-mtu`
`--rack-underlay-ebgp-peer-group`
`--rack-overlay-ebgp-peer-group`
`--lacp-timeout`
`--mct-port-channel`
`--control-vlan`
`--control-ve`
`--vni-auto-map`
`--backup-routing-enable`
`--backup-routing-ipv4-range`
`--backup-routing-ipv6-range`
	`--p2p-link-range`
	`--loopback-ip-range`
	`--bfd-enable`
	`--bfd-tx`
	`--bfd-rx`
	`--bfd-multiplier`
	`--mctlink-ip-range`
	`--single-rack-deployment`
	`--md5-password-enable`
	`--md5-password`
	`--bgp-dynamic-peer-listen-limit`
	`--bgp-multihop`
		`--spine-asn-block`
		`--super-spine-asn-block`
		`--leaf-asn-block`
		`--border-leaf-asn-block`
		`--configure-overlay-gateway`
		`--leaf-peer-group`
		`--spine-peer-group`
		`--super-spine-peer-group`
		`--optimized-replication-enable`
		`--mdtgroup-range`
		`--default-mdtgroup`

Fabric Settings in Active Fabric: Clos Fabric

Cannot be Updated	Can be Updated	Invalid Settings for Clos Fabric
--loopback-port-number --vtep-loopback-port-number --spine-asn-block --super-spine-asn-block --leaf-asn-block --border-leaf-asn-block --anycast-mac-address --ipv6-anycast-mac-address --mac-aging-timeout --mac-aging-conversation-timeout --mac-move-limit --duplicate-mac-timer --duplicate-mac-timer-max-count --configure-overlay-gateway --max-paths --allow-as-in --mtu --ip-mtu --leaf-peer-group --spine-peer-group --super-spine-peer-group --lacp-timeout --mct-port-channel --control-vlan --control-ve --vni-auto-map --optimized-replication-enable --backup-routing-enable --backup-routing-ipv4-range --backup-routing-ipv6-range --mdtgroup-range --default-mdtgroup	--p2p-link-range --loopback-ip-range --bfd-enable --bfd-tx --bfd-rx --bfd-multiplier --mctlink-ip-range --md5-password-enable --md5-password --bgp-dynamic-peer-listen-limit	--rack-l3-backup-ip-range --rack-asn-block --rack-border-leaf-asn-block --single-rack-deployment --rack-underlay-ebgp-peer-group --rack-overlay-ebgp-peer-group --bgp-multihop

Update md5-password on an Active Fabric

You can update an md5-password on an active fabric.

About This Task

When you update the md5-password on an active (already configured) fabric followed by the **efa fabric configure** operation, the **efa fabric configure** operation is considered successful even though the operational state (for example, CONN) of the fabric BGP peers (after **efa fabric configure**) is worse than the previous (before **efa**

fabric configure) operational state (for example, ESTABLISHED) of the fabric BGP peers. The system shows a warning message to indicate the worsened state of BGP peers.

Note
Run the **efa fabric topology show underlay --name <fabric-name>** command to get the latest status of the BGP session.

Procedure

1. Run the **efa fabric setting update** command.

```
efa fabric setting update --name fabric1 --md5-password-enable Yes --md5-password
'newpassword'
WARNING: configuring/clearing md5-password on an active fabric will result in BGP
neighbor sessions going down for a brief period when the fabric is reconfigured.
Please confirm if you want to continue with the fabric setting update [y/n]?y
fabric1 Fabric Update Successful
```

2. Run the **efa fabric configure** command.

```
efa fabric configure --name fabric1
Validate Fabric [Success]
Configure Fabric [Success]
10.25.225.11 : Operation[BGP Session(s) Clear Operation] has succeeded with the
warning:[BGP neighbor session 10.20.20.121 is in CONN state and could not be
established]
```

Update bgp-multihop on an Active Fabric

The bgp-multihop configuration is pushed for the non-Clos fabric only when the fabric contains multiple racks. The bgp-multihop configuration is applicable on the overlay ebgp peer-group.

About This Task

When you update the bgp-multihop on an active (already configured) fabric followed by the **efa fabric configure** operation, the **ebgp peer-group sessions** is reset. This is similar to the procedure followed during the update of md5 password followed by fabric configure.

Note
The fabric setting `bgp-multihop` is applicable for the non-Clos fabric and not applicable for the Clos fabric.

Procedure

1. Run the **efa fabric setting update** command.

```
efa fabric setting update --bgp-multihop 25 --name fab2
WARNING: configuring/clearing md5-password, configuring bgp-multihop on an active
fabric will result in BGP neighbor sessions going down for a brief period when the
fabric is reconfigured.
Please confirm if you want to continue with the fabric setting update [y/n]?
y
fab2 Fabric Update Successful
— Time Elapsed: 6.7189463s —
```

2. Run the **efa fabric show** command.

```
efa fabric show --name fab2
Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
settings-updated
Updated Fabric Settings: BGP-MULTIHOP
```

3. Run the **efa fabric configure** command.

```
efa fabric configure --name fab2
WARNING: 'fabric configure' will result in configuration change on the devices which
are in 'cfg refreshed' or 'fabric setting refreshed' or 'cfg refresh error' state.
Please check 'fabric show' to see the 'cfg refreshed' or 'fabric setting refreshed'
or 'cfg refresh error' devices. Please confirm if you want to continue with 'fabric
configure' [y/n]?
y
Validate Fabric [Success]
Configure Fabric [Success]
Please verify the fabric physical/underlay topology using 'efa fabric topology show
{physical | underlay}' before attempting tenant configuration on the fabric.
— Time Elapsed: 1m0.0693633s —
```

4. Complete the following configuration on SLX devices:

```
Freedom-06# show running-config router bgp
router bgp
local-as 4200000000
capability as4-enable
fast-external-fallover
neighbor overlay-ebgp-group peer-group
neighbor overlay-ebgp-group ebgp-multihop 25
neighbor underlay-ebgp-group peer-group
```

Example

- The following is an example of **efa fabric show** command output:

```
efa fabric show --name fab2
Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
configure-success
-------------------------------------------------------------------------------------
-----------------------
IP ADDRESS    RACK    HOST NAME    ASN     ROLE  DEVICE STATE  APP STATE   CONFIG
PENDING  VTLB LB
                                                                                  GEN

REASON CONFIGS  ID   ID
10.20.246.24  room1-rack2  Freedom-08  4200000001  leaf  provisioned  cfg in-sync
NA      NA       2   1
10.20.246.23  room1-rack2  Freedom-07  4200000001  leaf  provisioned  cfg in-sync
NA      NA       2   1
10.20.246.21  room1-rack1  Freedom-05  4200000000  leaf  provisioned  cfg in-sync
NA      NA       2   1
10.20.246.22  room1-rack1  Freedom-06  4200000000  leaf  provisioned  cfg in-sync
NA      NA       2   1
```

- The following is an example of SLX configuration **before** fabric setting update and fabric configure:

```
Freedom-06# show running-config router bgp
router bgp
local-as 4200000000
capability as4-enable
fast-external-fallover
neighbor overlay-ebgp-group peer-group
neighbor overlay-ebgp-group ebgp-multihop 4
```

- The following is an example of overlay EBGP session states **before** the fabric configure:

```
Neighbor Address AS#        State Time     Rt:Accepted Filtered Sent To Send
172.31.254.107   4200000001 ESTAB 0h22m47s 0           0        0    0
172.31.254.116   4200000001 ESTAB 0h23m9s  0           0        0    0
```

- The following is an example of non-Overlay EBGP Session states before the fabric configure:

```
Neighbor Address AS#        State Time     Rt:Accepted Filtered Sent To Send
10.10.10.1       4200000001 ESTAB 0h23m43s 3           0        3    0
10.20.20.6       4200000000 ESTAB 0h54m25s 2           0        5    0
```

- The following is an example of overlay EBGP Session states **after** the fabric configure:

```
Neighbor Address AS#        State Time    Rt:Accepted Filtered Sent To Send
172.31.254.107   4200000001 ESTAB 0h0m24s 0           0        0    0
172.31.254.116   4200000001 ESTAB 0h0m47s 0           0        0    0
```

- The following is an example of non-Overlay EBGP Session states after the fabric configure:

```
Neighbor Address AS#        State Time     Rt:Accepted Filtered Sent To Send
10.10.10.1       4200000001 ESTAB 0h26m13s 3           0        3    0
10.20.20.6       4200000000 ESTAB 0h56m55s 2           0        5    0
```

Fabric Configuration using Force

You can configure fabric using force only via CLI.

The fabric configuration with `force` option from the REST API is deprecated from XCO 3.2.0 onwards. It is supported only via CLI.

```
http://localhost:8081/v1/fabric/configure?fabric-name=fabg&force=true
Response: 500 Internal server
[
{
"ip-address": "",
"error": [
{ "code": 1727, "message": "force operation is not supported using REST API for this
release use CLI" }
]
}
]
```

Fabric DRC using Force

If the fabric device is in the `cfg-in-sync` state, the **`efa inventory drift-reconcile execute --ip <device-ip>`** command does not result in fabric drift identification and reconciliation.

To force drift identification and reconciliation of the fabric configuration on a fabric device, use the **`efa fabric debug device drift --device-ip <device-ip> --name <fabric-name> --force`** command.

Fabric Event Handling

Event handling reason code is generated only in case of a drift in configuration at the device after XCO fabric is configured. The following listed events generated by RASlog event or device update is handled in fabric services:

The **efa fabric show** command displays all the events handled at fabric services.

```
CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable
```

Drift (SLX vs XCO)	App State	CONFIG GEN REASON (efa fabric show)
Evpn	Cfg-refresh	EU/ED - Evpn Delete or Update
Cluster	Cfg-refresh	MD/MU - MCT Delete or Update
Router bgp	Cfg-refresh	BGPU - Router BGP Update
Overlay gateway	Cfg-refresh	OD/OU - Overlay Gateway Delete/ Update,
Interface (Phy/ Lo/Po)	Cfg-refresh	IA/ID/IU - Interface Add, Delete or Update
LLDP	Cfg-refresh	LA/LD - Link Add or Delete

Importing a Fabric Database

You can import the fabric databases (including devices) of EFA versions (EFA 2.5.5 and later). The support of an older database by newer software is often called Brownfield support.

XCO supports the import (or migration) of fabric databases of earlier versions of EFA (2.5.5 and later) in the following ways:

- You can import devices that were configured in the fabric of the earlier version of EFA (2.5.5 and later).
- You can import the fabric that was configured in the earlier version of EFA (2.5.5 and later).
- You can import and configure the fabric and the devices using the **efa fabric import** and **efa fabric configure** commands.
- You can import only 3-stage Clos fabrics because only 3-stage Clos fabrics were supported in earlier versions of EFA (2.5.5 and later).

When you run the **`efa fabric import`** command, XCO performs the following tasks:

- Learns (fetches) the device configuration and registers the devices with the XCO Inventory service. If registration fails for any device, XCO removes all devices from the inventory. You must fix the errors and rerun the command.
- Fetches the older fabric settings and creates a new fabric. You provide the new fabric name when you run the command. If the fabric name you provide already exists, with devices added to it, then an error is returned. You must fix the error and then rerun the command.
- Adds the registered devices to the new fabric and validates the configurations for global devices, interfaces, MCT, overlay gateway, EVPN, and BGP. For any conflict in the configuration, XCO deletes the devices from the fabric, unregisters the devices, and deletes the new fabric. The devices retain their pre-import configuration. You must fix the configuration conflicts and then rerun the command.

Note
The following configuration deviations are allowed: the MCT cluster name, the EVPN name, the overlay gateway name, and the MCT cluster node ID. In each case, the value is taken from the device.

Upon successful migration, devices that are added to the fabric have the provisioning state of `CFG-READY`. You then run the **`efa fabric configure`** command to transition the devices to the `CFG-IN-SYNC` state.

For more information, see #unique_115.

Pre-validation of Configuration

When devices with preexisting configuration are added to XCO, fabric service performs validations before adding devices to the XCO fabric.

If any of the configuration that is retrieved from the devices does not fall under the fabric settings range, an error is displayed. You can perform corrective actions to add such devices to the XCO fabric. For more information, see the following tables.

Global Device Configuration

Use Case	Valid	Invalid
L2 MTU	If the value fetched from the device is same as that configured in fabric settings.	If the value does not match, an error is displayed.
IP MTU	If the value fetched from the device is same as that configured in fabric settings.	If the value does not match, an error is displayed.
ASN	If the value fetched from the device is within the ASN range configured in fabric settings.	If the value is out of range of what is configured in fabric settings. If the value is conflicting with existing device which is already added to Fabric.

Interface Configuration

Use Case	Valid	Invalid
Interface IP Address	Received an IP address within “Link IP range” of fabric settings.	• If IP address received is out of range, a validation error is displayed. • If IP address received is within the range but is already in use/reserved by fabric, a validation error is displayed.
Loopback Interface ID	Loopback Interface ID is reconciled	
VTEP Loopback Interface ID	VTEP Loopback interface ID is reconciled	
Loopback Interface IP Address	Loopback IP address is within “Loopback IP range” of fabric settings.	If the loopback IP address is out of range of “Loopback IP range” of fabric settings, an error is displayed.
VE IP address	VE IP address is within “MCT Link IP Range” of fabric settings.	• If IP address received is out of range, a validation error is displayed. • If IP address received is within the range but is already in use/reserved by fabric, a validation error is displayed.
Static IP Route (Applicable for SLX 9540 and SLX 9640)	If the IP route is same as fabric intended configuration.	If the nexthop does not point to VE IP.

MCT Configuration

MCT peer and VE validations are platform specific.

Use Case	Valid	Invalid
MCT Cluster Name	MCT Cluster Name is different from fabric name in fabric properties. MCT Cluster Name learned from device is used while configuring the device, so that the cluster name is reconciled.	
Cluster Control VLAN	Cluster Control VLAN matches with “Control VLAN” of fabric settings.	If the VLAN does not match, an error is displayed.
Cluster Control VE	Cluster Control VE matches with “Control VE” of fabric settings	If the VE does not match, an error is displayed.
MCT Peer IP	Peer IP address is within IP range of “MCT Link IP range” of fabric settings.	If Peer IP address is out of IP range, an error is displayed.
MCT Peer Interface	Peer Interface type matches with fabric settings and ID is reconciled.	

Overlay Gateway Configuration

Use Case	Valid	Invalid
Overlay Gateway Name	Gateway Name is reconciled.	
VNI Auto Map	VNI auto mapping setting configured on the device matches with fabric settings.	If gateway is in activated state and VNI Auto Map setting is different from the fabric settings, an error is displayed.
Overlay Gateway Interface	Gateway Interface, for example loopback 2 is reconciled.	

EVPN Configuration

Use Case	Valid	Invalid
EVPN Name	EVPN Name is reconciled	
MAC Aging Timeout (check based on device capability, applies to SLX 9140 and SLX 9240)	Field value is overwritten by the fabric settings value.	

Use Case	Valid	Invalid
MAC Aging Conversation Timeout (check based on device capability, applies to SLX 9140 and SLX 9240)		
MAC Move Limit (check based on device capability)		
ArpAgingTimeout (check based on device capability)		
Duplicate Mac Timer		
Duplicate MAC Timer MAX Count		

BGP Configuration

To pre-validate the BGP configuration, the BGP configuration must be prepared similar to the add device phase. Once the BGP configuration is computed, the configuration retrieved from the device is compared against it.

Use Case	Valid	Invalid
Router ID	If the generated router id matches the one received from the device.	If the router id does not match, an error is displayed.
BFD Enable/Disable	If the BFD value from device matches the one that is computed from fabric settings. While configuring the fabric, the values computed by fabric service override the ones on the device.	NA
BFD Tx/Rx Timer Values	If the values from device match with the ones that are generated. While configuring the fabric, the values computed by fabric service override the ones on the device.	NA
Network Address	If the value is within “Loopback IP Range” of fabric settings or matches the computed value.	If the value is out of range or clashes with another IP neighbor already stored in fabric DB, an error is displayed.

Use Case	Valid	Invalid
EVPN Neighbor IP Address	If the neighbor IP address falls in range of “Link IP range” of fabric settings. There may be a case where the neighbor IP address is valid but neighbor is not part of fabric. You can ignore such validation as that configuration is a no-op for fabric.	If the value is out of range or clashes with another IP neighbor already stored in fabric DB, an error is displayed.
Remote ASN	If the ASN is within the range of fabric settings and not already in use by another neighbor.	If the ASN is out of range or already reserved.
Peer group name	If the peer group name matches the peer group that is computed.	If the value does not match, an error is displayed.

BGP Tables

The following BGP tables help in computing the diffs for the events from the inventory service.

- Router BGP table
- BGP peer group table
- BGP IP address family table
- BGP IP neighbor address table
- BGP EVPN address family table
- BGP EVPN neighbor address table

All BGP tables handle the DB migration so that upgrade from older XCO to newer XCO works.

For more information on the attributes of each table, refer to Database schema section or `fabric_schema.sql` file.

BGP Events

The following BGP events from inventory service are handled as part of event handling.

BGP Router Delete

When router BGP delete message is received, fabric passes through all the IP and EVPN neighbors, peer group tables and the entries corresponding to the device for which router BGP delete message is received and mark the entries as ‘create’ to configure the router BGP and its related neighbors on the device.

BGP IP Neighbor Delete

When BGP IP neighbor delete message is received, fabric passes through all the IP neighbors deleted and which exists in fabric database for a given neighbor IP or remote ASN and mark the entries as ‘create’ to configure the deleted IP neighbors on the device.

BGP IP Neighbor Update

When BGP IP neighbor update message is received, fabric passes through all the IP neighbors matching the neighbor IP for the device in the database. If any of the fabric managed attribute in the IP neighbor table is changed, fabric marks the entries as ‘update’ and pushes the configuration back to the device.

BGP EVPN Neighbor Delete

When BGP EVPN neighbor delete message is received, fabric passes through all the EVPN neighbors deleted and which exists in fabric database for a given neighbor IP or remote ASN and mark the entries as ‘create’ to configure back the deleted EVPN neighbors on the device.

BGP EVPN Neighbor Update

When BGP EVPN neighbor update message is received, fabric passes through all the EVPN neighbors matching the neighbor IP for the device in database. If any of the fabric managed attribute in the EVPN neighbor table is changed, fabric marks the entries as ‘update’ and pushes the configuration back to the device.

Peer Group Delete

Peer Group Delete message is received only when there are no IP/EVPN neighbors associated with it. If there are no IP/EVPN neighbors associated with it, fabric marks the Peer Group as ‘delete’.

Peer Group Update

When Peer group attributes such as BFD and remote ASN change, inventory sends a peer group update message. The fabric processes this message and checks if the peer group exists in the database. If the peer group exists and there are changes to the attributes, the fabric pushes the peer group configuration with fabric intended configuration back to the device.

BGP IP Address Family Delete

BGP IP Address Family Delete message is received when the IP address-family for a device is deleted through CLI or out-of-band means. When fabric receives this message, it passes through all the IP neighbors associated with that address-family and marks the entries as ‘create config’ to restore all the deleted IP neighbors on the device.

BGP EVPN Address Family Delete

BGP EVPN Address Family Delete message is received when the EVPN address-family for a device is deleted through CLI or out-of-band means. When fabric receives the message, it passes through all the EVPN neighbors associated with that address-family

and marks the entries as ‘create config’ to restore all the deleted EVPN neighbors on the device.

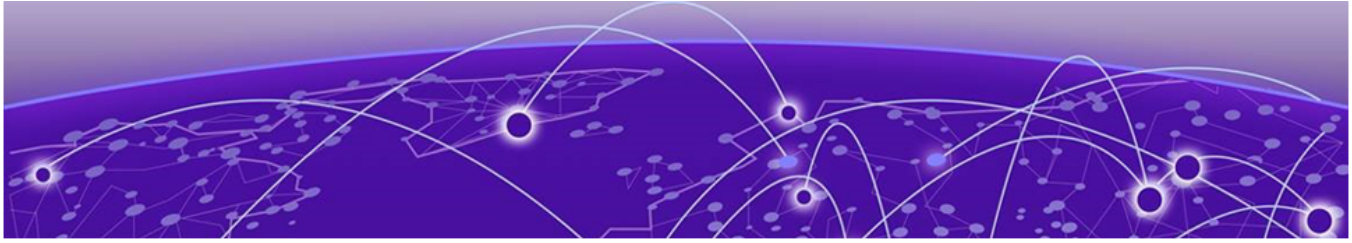

Tenant Service Provisioning

Tenant Services Provisioning Overview

Tenant Services exposes the CLI and REST API for automating the Tenant network configuration on the Clos and non-Clos (small data center) overlay fabric.

Tenant network configuration includes VLAN, BD, VE, EVPN, VTEP, VRF, and Router BGP configuration on the necessary fabric devices to provide L2-Extension, L3-Extension across the fabric, L2-Handoff, and L3-Handoff at the edge of the fabric.

Tenant Services provisioning automates the Tenant configuration, which can be a subset of the combinations provided by the switching hardware.

Tenant Services supports multiple fabrics.

Note

You cannot perform fabric and tenant operations when manual DRC is in progress.

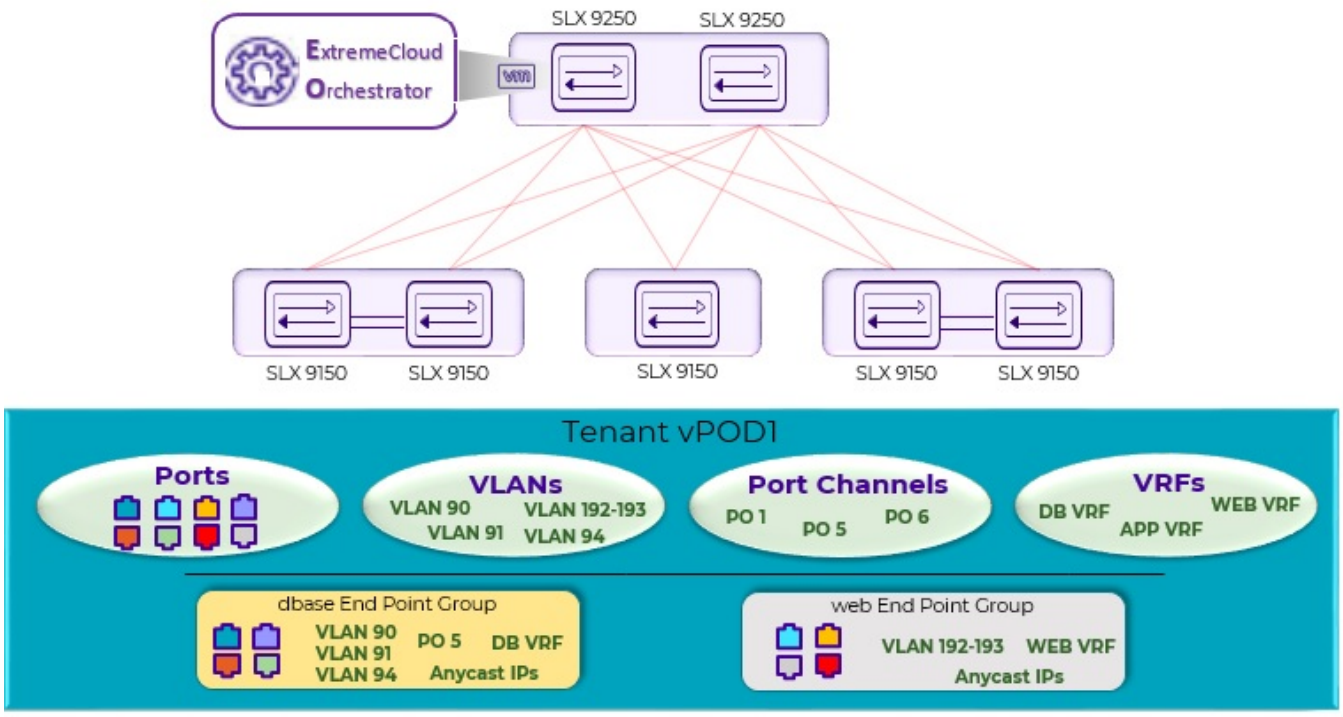

Figure 15: Tenant Services Overview

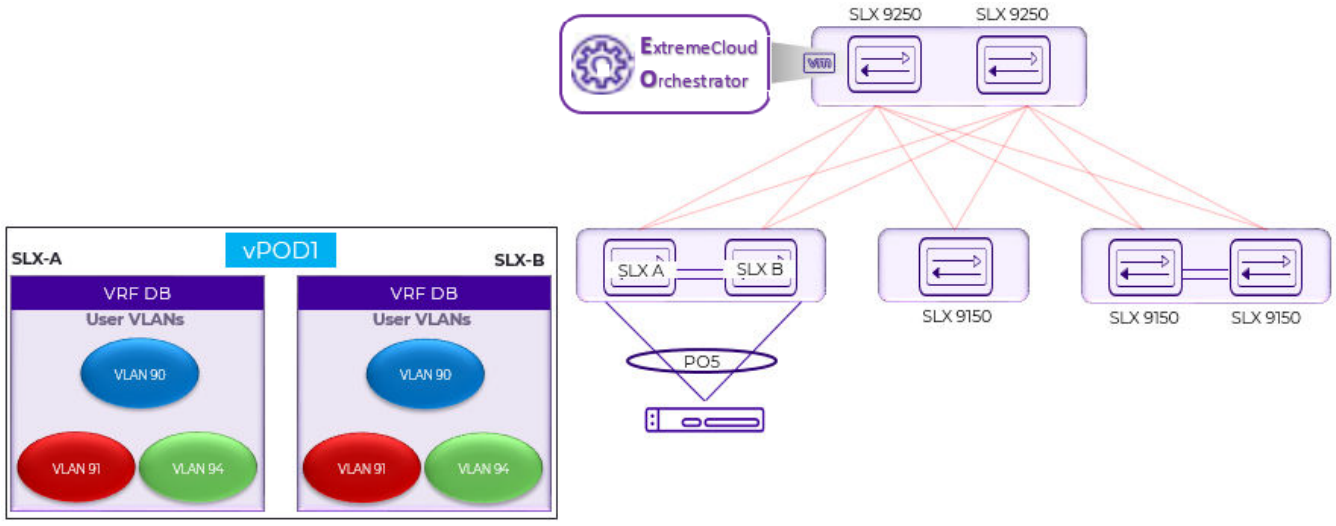

Figure 16: Tenant Name vPOD1, VRF Name DB

Tenant

A Tenant is a logical construct that owns resources as follows:

- VLAN range: Ctags pertaining to which the traffic is expected to ingress and egress.

  Note
  Customer VLAN tag (Ctag) is used to identify the customer broadcast domain. In the IP fabric network, it represents the customer and is mapped into a VXLAN tunnel through a virtual network identifier (VNI). The VNI is the ID used to identify the VXLAN tunnel. With auto VNI mapping, the Ctag ID equals the VNI. You can also manually map Ctags to user-defined VNIs. These VNIs can be mapped to VLAN IDs (up to 4k) or to BD's (bridge domain) IDs.

- Device ports: Ports on which the traffic is expected to ingress and egress.

VLAN-based Tenant

For a VLAN-based tenant, realization of network on the device is done using VLAN and switchport VLANs. Bridge domains are used for EVPN IRB.

Bridge domain-based Tenant

For a BD-based tenant, realization of network on the device is done using BD and BD-LIF. BD is used for EVPN IRB.

Scalability

Table 14: VNI scalability

VNI type	Scale
Non-auto VNI mapping	• The number of VNI (networks) supported per device = 8K [4K VLAN + 4K BD] • The maximum number of VNI (networks) supported in the fabric = [8K * number of devices in the fabric].
Auto VNI mapping	• The number of VNI (networks) supported per device = 8K [4K VLAN + 4K BD] • The number of VNI (networks) supported per fabric = 8K

Event handling

Event handling specifies the scope of the tenant configuration on the devices.

Devices are added to the Tenant service only when the fabric is provisioned on the devices.

An event is an occurrence of a device being removed from the fabric or from the Inventory.

- When a device is removed from the fabric or inventory, the device is cleaned up from Tenant Service and the Tenant configuration is removed from the device.
- User-created entities, such as Tenant, VRF, and Endpoint Group, are not deleted whereas references for ports or port-channels of deleted devices are removed.

Clos Fabric with Non-auto VNI Maps

Auto VNI simplifies the mapping IDs by using the VLAN ID as the VNI ID, for example VLAN 100 = VNI 100.

This method of mapping works well in environments where overlapping VLANs are not being used. However, if two different tenants are using VLAN 100, VNI 100 cannot be used by both. At this point, manual mapping of VLAN to VNI is required. ExtremeCloud Orchestrator simplifies this process by allowing VNI ranges for tenants to automate "manual" mapping to work for overlapping VLANs.

The following figure shows a 3-stage Clos topology.

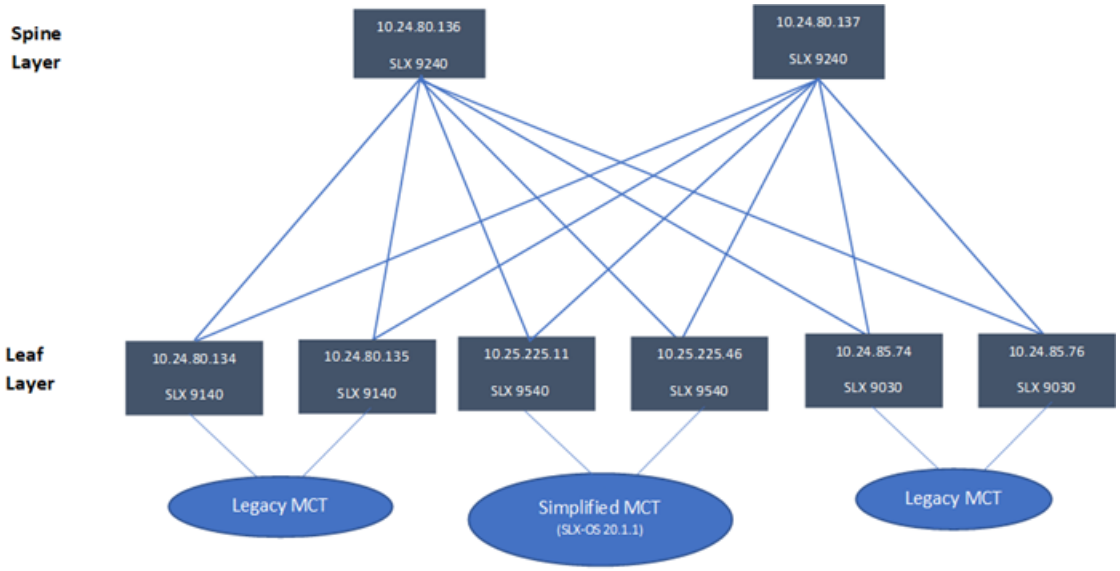

Figure 17: 3-stage Clos topology

The following commands configure the 3-stage Clos topology:

```
efa fabric create --name fabric1

efa fabric setting update --name fabric1 --vni-auto-map No

efa fabric device add-bulk --spine 10.24.80.136 --border-leaf 10.25.225.11,10.25.225.46
--leaf 10.24.80.134-135,10.24.85.74,10.24.85.76 --username admin --password password
--name fabric1

efa fabric configure --name fabric1
```

The following figure shows tenant constructs in the Clos Fabric.

Figure 18: Scope of tenant constructs

Clos Fabric with Auto VNI Map

- In Clos fabric with auto VNI map, the VNI is statically derived using the VLAN ID or BD ID.
  - For the VLAN case, VNI = VLAN ID
  - For the BD case, VNI = 4096 + BD ID
  - You cannot reserve l2-vni-range or l3-vni-range for a given tenant.
  - You cannot provide a specific l2-vni or l3-vni in an endpoint group.
- VLAN Based Tenants:

  Multiple VLAN based tenants cannot share the same VLAN, considering the multiple tenants cannot share the same VNI.
- BD Based Tenants:

  Multiple BD based tenants can share the same VLAN, as the VLANs from each tenant are mapped to a unique BD and further a unique VNI.

Multi Tenancy

XCO supports multi tenancy by allowing multiple tenants to have overlapping ctags and non-overlapping L2VNI. A tenant ctag will get a unique L2VNI and a unique network allocated in the fabric

The following example shows a multi tenancy configuration.

```
efa tenant create --name tenant11 --vrf-count 10 --vlan-range 2-4090 --port
10.24.80.134[0/15-17],10.24.80.135[0/15-17],10.25.225.11[0/15-17],10.25.225.46[0/15-17],10
.24.85.74[0/15-17],10.24.85.76[0/15-17] --description Subscriber1

efa tenant show
+----------+--------+--------+-------+-------+-------+-----------------------+
|   Name   | L2VNI  | L3VNI  | VLAN  | VRF   | Enable|         Ports         |
|          | -Range | -Range | -Range| -Count| -BD   |                       |
+----------+--------+--------+-------+-------+-------+-----------------------+
| tenant11 |        |        | 2-4090| 10    | False | 10.24.85.74[0/15-17]  |
|          |        |        |       |       |       | 10.24.80.135[0/15-17] |
|          |        |        |       |       |       | 10.25.225.11[0/15-17] |
|          |        |        |       |       |       | 10.25.225.46[0/15-17] |
|          |        |        |       |       |       | 10.24.80.134[0/15-17] |
|          |        |        |       |       |       | 10.24.85.76[0/15-17]  |
+----------+--------+--------+-------+-------+-------+-----------------------+

efa tenant create --name tenant12 --vrf-count 10  --vlan-range 2-4090 --port
```

```
10.24.80.134[0/18-20],10.24.80.135[0/18-20],10.25.225.11[0/18-20],10.25.225.46[0/18-20],10
.24.85.74[0/18-20],10.24.85.76[0/18-20]
Tenant Creation Failed:
         Vlan (2) overlaps with Tenant (tenant11)

efa tenant create --name tenant21 --vrf-count 10 --enable-bd --port
10.24.80.134[0/21-25],10.24.80.135[0/21-25],10.24.85.74[0/21-25],10.24.85.76[0/21-25],10.2
5.225.11[0/21-25],10.25.225.46[0/21-25]

efa tenant create --name tenant22 --vrf-count 10 --enable-bd --port
10.24.80.134[0/26-30],10.24.80.135[0/26-30],10.24.85.74[0/26-30],10.24.85.76[0/26-30],10.2
5.225.11[0/26-30],10.25.225.46[0/26-30]

efa tenant show
+----------+-------+-------+-------+-------+-------+----------------------+
|   Name   | L2VNI | L3VNI | VLAN  | VRF   | Enable|         Ports        |
|          | -Range| -Range| -Range| -Count| -BD   |                      |
+----------+-------+-------+-------+-------+-------+----------------------+
| tenant11 |       |       | 2-4090| 10    | False | 10.25.225.46[0/15-17] |
|          |       |       |       |       |       | 10.25.225.11[0/15-17] |
|          |       |       |       |       |       | 10.24.80.135[0/15-17] |
|          |       |       |       |       |       | 10.24.85.74[0/15-17]  |
|          |       |       |       |       |       | 10.24.85.76[0/15-17]  |
|          |       |       |       |       |       | 10.24.80.134[0/15-17] |
+----------+-------+-------+-------+-------+-------+----------------------+
| tenant21 |       |       | 2-4090| 10    | True  | 10.24.85.74[0/21-25]  |
|          |       |       |       |       |       | 10.25.225.11[0/21-25] |
|          |       |       |       |       |       | 10.25.225.46[0/21-25] |
|          |       |       |       |       |       | 10.24.80.134[0/21-25] |
|          |       |       |       |       |       | 10.24.85.76[0/21-25]  |
|          |       |       |       |       |       | 10.24.80.135[0/21-25] |
+----------+-------+-------+-------+-------+-------+----------------------+
| tenant22 |       |       | 2-4090| 10    | True  | 10.24.85.76[0/26-30]  |
|          |       |       |       |       |       | 10.25.225.11[0/26-30] |
|          |       |       |       |       |       | 10.24.80.135[0/26-30] |
|          |       |       |       |       |       | 10.25.225.46[0/26-30] |
|          |       |       |       |       |       | 10.24.85.74[0/26-30]  |
|          |       |       |       |       |       | 10.24.80.134[0/26-30] |
+----------+-------+-------+-------+-------+-------+----------------------+

efa tenant epg create --name epg11 --tenant tenant11 --po po1115,po1215,po1315 --
switchport-mode trunk --switchport-native-vlan 11 --ctag-range 11-12

efa tenant epg create --name epg21 --tenant tenant21 --po po2121,po2221,po2321 --
switchport-mode trunk  --ctag-range 11-12

efa tenant epg create --name epg22 --tenant tenant21 --po po2122,po2222,po2322  --
switchport-mode trunk  --ctag-range 11-12

efa tenant epg show

=========================================================================
Name          : epg11
Tenant        : tenant11
Description   :
Ports         :
POs           : po1315, po1215, po1115
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12

Network Property [Flags : * - Native Vlan]
+------+--------+------------+---------+-------------+-------------+
| Ctag | L2-Vni | Anycast-ip | BD-name |  Dev-state  |  App-state  |
```

```
+------+--------+------------+---------+-------------+-------------+
| 11*  | 11     |            |         | provisioned | cfg-in-sync |
+------+--------+------------+---------+-------------+-------------+
| 12   | 12     |            |         | provisioned | cfg-in-sync |
+------+--------+------------+---------+-------------+-------------+

==========================================================================
==========================================================================
Name          : epg21
Tenant        : tenant21
Description   :
Ports         :
POs           : po2121, po2221, po2321
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12

Network Property [Flags : * - Native Vlan]
+------+--------+------------+--------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-ip |   BD-name    |  Dev-state  |   App-state |
+------+--------+------------+--------------+-------------+-------------+
| 11   | 4099   |            | Auto-BD-4099 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+
| 12   | 4100   |            | Auto-BD-4100 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+

==========================================================================
==========================================================================
Name          : epg22
Tenant        : tenant21
Description   :
Ports         :
POs           : po2122, po2222, po2322
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12

Network Property [Flags : * - Native Vlan]
+------+--------+------------+--------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-ip |   BD-name    |  Dev-state  |  App-state  |
+------+--------+------------+--------------+-------------+-------------+
| 12   | 4102   |            | Auto-BD-4102 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+
| 11   | 4101   |            | Auto-BD-4101 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+

efa tenant epg create --name epg23 --tenant tenant21 --po po2122,po2322  --switchport-
mode trunk  --ctag-range 21-22 --bridge-domain 21:Auto-BD-4101 --bridge-domain 22:Auto-
BD-4102

efa tenant epg show

==========================================================================
Name          : epg11
Tenant        : tenant11
Description   :
Ports         :
POs           : po1315, po1215, po1115
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12

Network Property [Flags : * - Native Vlan]
+------+--------+------------+---------+-------------+-------------+
```

```
| Ctag | L2-Vni | Anycast-ip | BD-name |  Dev-state  |  App-state  |
+------+--------+------------+---------+-------------+-------------+
| 11*  | 11     |            |         | provisioned | cfg-in-sync |
+------+--------+------------+---------+-------------+-------------+
| 12   | 12     |            |         | provisioned | cfg-in-sync |
+------+--------+------------+---------+-------------+-------------+

==========================================================================
==========================================================================
Name          : epg21
Tenant        : tenant21
Description   :
Ports         :
POs           : po2121, po2221, po2321
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12

Network Property [Flags : * - Native Vlan]
+------+--------+------------+--------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-ip |   BD-name    |  Dev-state  |   App-state |
+------+--------+------------+--------------+-------------+-------------+
| 11   | 4099   |            | Auto-BD-4099 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+
| 12   | 4100   |            | Auto-BD-4100 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+

==========================================================================
==========================================================================
Name          : epg22
Tenant        : tenant21
Description   :
Ports         :
POs           : po2122, po2222, po2322
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12

Network Property [Flags : * - Native Vlan]
+------+--------+------------+--------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-ip |   BD-name    |  Dev-state  |  App-state  |
+------+--------+------------+--------------+-------------+-------------+
| 12   | 4102   |            | Auto-BD-4102 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+
| 11   | 4101   |            | Auto-BD-4101 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+

==========================================================================
==========================================================================
Name          : epg23
Tenant        : tenant21
Description   :
Ports         :
POs           :
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 21-22

Network Property [Flags : * - Native Vlan]
+------+--------+------------+--------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-ip |   BD-name    |  Dev-state  |  App-state  |
+------+--------+------------+--------------+-------------+-------------+
| 21   | 4101   |            | Auto-BD-4101 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+
```

```
| 22   | 4102   |            | Auto-BD-4102 | provisioned | cfg-in-sync |
+------+--------+------------+--------------+-------------+-------------+

=========================================================================
```

Provision a Tenant Entity

A tenant is a group of users that own or have access to shared resources.

About This Task

Complete the following tasks to provision a tenant in your XCO fabric.

Procedure

1. Create a Tenant on page 146
2. Update a Tenant on page 147
3. Show a Tenant on page 149
4. Delete a Tenant on page 151
5. Configure a Tenant on page 151
6. Share Resources Across Tenants using Shared Tenant on page 344

Create a Tenant

You can specify the resources like device ports, VLAN range, L2 VNI range, L3 VNI range, and VRF count when you create a tenant.

About This Task

Follow this procedure to set up a logical construct called tenant.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

Run the **`efa tenant create`** command.

```
efa tenant create --name <tenant-name> --description <tenant-description> --l2-vni-range
<value> --l3-vni-range <value> -- vlan-range <value> --vrf-count <value> --enable-bd –
port <list-of-ports>
```

Example

- The following example shows how a VLAN-based tenant is created with manual VNI mapping:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant11 --l2-vni-range
10002-14190
--l3-vni-range 14191-14200 --vlan-range 2-4090 --vrf-count 10 --port
10.20.216.15[0/11-20],10.20.216.16[0/11-20]
--description Subscriber1

Tenant created successfully.
```

```
--- Time Elapsed: 455.141597ms ---
```

- The following example shows how a BD-based tenant is created with manual VNI mapping:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant21 --l2-vni-range
30002-34190
--l3-vni-range 34191-34200 --vlan-range 2-4090 --vrf-count 10 --enable-bd
--port 10.20.216.15[0/21-28],10.20.216.16[0/21-28]

Tenant created successfully.

--- Time Elapsed: 501.176996ms ---
```

- The following example shows how a tenant is created with auto-VNI mapping and breakout ports:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant12
--vlan-range 2-100 --vrf-count 10 --port
10.20.216.103[0/1-10],10.20.216.104[0/1-5,0/6:1-4]

Tenant created successfully.

--- Time Elapsed: 427.73527ms ---
```

- The following example shows how a shared tenant is created with shared ports:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name ST
--type shared --port 10.20.216.15[0/1-10],10.20.216.16[0/1-10]

Tenant created successfully.

--- Time Elapsed: 381.182892ms ---
```

Update a Tenant

You can update the tenant attributes using the operations, such as desc-update, vni-update, port-add, port-delete, vlan-add, vlan-delete, vlan-update, num-vrf-update, and enable-bd-update.

About This Task

Follow this procedure to update an existing tenant.

Procedure

Run the **`efa tenant update`** comamnd.

```
efa tenant update --name <tenant-name> --operation <value> --description <tenant-
description> --l2-vni-range <value> --l3-vni-range <value> --vlan-range <value> --vrf-
count <value> --enable-bd –port <list-of-ports> --force
```

Example

The following example shows the existing tenant configuration:

```
(efa:root)root@node-2:~# efa tenant show --detail
============================================================
Name            :  tenant11
Type            :  private
Description     :
VLAN Range      :  2001-2150
L2VNI Range     :
L3VNI Range     :
```

```
VRF Count       :  100
Enable BD       :  false
Ports           :  10.20.246.6[0/1-10]
                :  10.20.246.5[0/1-10]

============================================================

--- Time Elapsed: 159.855317ms ---
```

- The following example updates existing tenant description:

```
(efa:root)root@node-2:~# efa tenant update --name tenant11 --operation desc-update --
description tenant11Desc

Tenant updated successfully.

--- Time Elapsed: 109.837946ms ---
```

- The following example updates existing tenant description. This operation is allowed only when no EPG is associated with tenant:

```
(efa:root)root@node-2:~# efa tenant update --name tenant11 --operation enable-bd-
update --enable-bd

Tenant updated successfully.

--- Time Elapsed: 92.004637ms ---
```

- The following example updates the L2 VNI and L3 VNI for an existing tenant:

```
(efa:root)root@node-2:~# efa tenant update --name tenant11 --operation vni-update --l2-
vni-range 10002-14190 --l3-vni-range 14191-16200

Tenant updated successfully.

--- Time Elapsed: 97.955561ms ---
```

- The following example updates the VLAN for an existing tenant:

```
(efa:root)root@node-2:~# efa tenant update --name tenant11 --operation vlan-update --
vlan-range 2-4090

Tenant updated successfully.

--- Time Elapsed: 138.503235ms ---
```

- The following example updates the L2 VNI and L3 VNI for an existing tenant:

```
(efa:root)root@node-2:~# efa tenant update --name tenant11 --operation num-vrf-update
--vrf-count 10

Tenant updated successfully.

--- Time Elapsed: 93.855235ms ---
```

- The following example updates the ports for an existing tenant:

```
(efa:root)root@node-2:~# efa tenant update --name tenant11 --operation port-add --port
10.20.246.5[0/15-17],10.20.246.6[0/15-17]

Tenant updated successfully.

--- Time Elapsed: 185.372609ms ---


(efa:root)root@node-2:~# efa tenant show --detail
================================================
Name            :  tenant11
Type            :  private
Description     :  tenant11Desc
```

```
VLAN Range      :  2-4090
L2VNI Range     :  10002-14190
L3VNI Range     :  14191-16200
VRF Count       :  10
Enable BD       :  true
Ports           :  10.20.246.5[0/1-10,0/15-17]
                :  10.20.246.6[0/1-10,0/15-17]

================================================

--- Time Elapsed: 69.527552ms ---
```

Show a Tenant

You can view a brief or detailed output of all the tenants or a given tenant.

About This Task

Follow this procedure to show a tenant configuration.

Procedure

Run the **`efa tenant show`** command.

Example

- The following example shows brief output of all VRFs:

```
(efa:root)root@node-2:~# efa tenant show
+----------+--------+-------+-------------+-------------+-----+-------
+-----------------------------+
|   Name   |  Type  | VLAN  | L2VNI
Range | L3VNI Range |VRF  | Enable|            Ports            |
|          |        | Range |             |
|Count| BD    |                             |
+----------+--------+-------+-------------+-------------+-----+-------
+-----------------------------+
| tenant11 | private| 2-4090| 10002-14190
| 14191-16200 |10   | true  | 10.20.246.5[0/1-10,0/15-17] |
|          |        |       |             |
|     |       | 10.20.246.6[0/1-10,0/15-17] |
+----------+--------+-------+-------------+-------------+-----+-------
+-----------------------------+
| tenant22 | shared |  2-21 |
|             |100  | false |      10.20.246.5[0/11]      |
|          |        |       |             |
|     |       |      10.20.246.6[0/11]      |
+----------+--------+-------+-------------+-------------+-----+-------
+-----------------------------+
Tenant Details

--- Time Elapsed: 163.599377ms ---
```

- The following example shows detailed output of all tenants:

```
(efa:root)root@node-2:~# efa tenant show --detail
============================================================
Name            :  tenant11
Type            :  private
Description     :  tenant11Desc
VLAN Range      :  2-4090
L2VNI Range     :  10002-14190
L3VNI Range     :  14191-16200
VRF Count       :  10
Enable BD       :  true
Ports           :  10.20.246.6[0/1-10,0/15-17]
```

```
                :  10.20.246.5[0/1-10,0/15-17]

============================================================
============================================================
Name            :  tenant22
Type            :  shared
Description     :
VLAN Range      :  2-21
L2VNI Range     :
L3VNI Range     :
VRF Count       :  100
Enable BD       :  false
Ports           :  10.20.246.5[0/11]
                :  10.20.246.6[0/11]

============================================================

--- Time Elapsed: 145.207842ms ---
```

- The following example shows brief output of a specific tenant:

```
(efa:root)root@node-2:~# efa tenant show --name tenant11
+----------+---------+-------+-------------+------------+------+-------
+------------------------------+
|   Name   |  Type   | VLAN  | L2VNI
Range |L3VNI Range | VRF  | Enable|            Ports            |
|          |         | Range |             |
| Count| BD    |                             |
+----------+---------+-------+-------------+------------+------+-------
+------------------------------+
| tenant11 | private | 2-4090| 10002-14190
|14191-16200 | 10   | true  | 10.20.246.5[0/1-10,0/15-17] |
|          |         |       |             |
|      |       | 10.20.246.6[0/1-10,0/15-17] |
+----------+---------+-------+-------------+------------+------+-------
+------------------------------+
Tenant Details

--- Time Elapsed: 51.249187ms ---
```

- The following example shows detailed output of specific tenant:

```
(efa:root)root@node-2:~# efa tenant show --name tenant11 --detail
============================================================
Name            :  tenant11
Type            :  private
Description     :  tenant11Desc
VLAN Range      :  2-4090
L2VNI Range     :  10002-14190
L3VNI Range     :  14191-16200
VRF Count       :  10
Enable BD       :  true
Ports           :  10.20.246.5[0/1-10,0/15-17]
                :  10.20.246.6[0/1-10,0/15-17]
============================================================
--- Time Elapsed: 79.930076ms ---
```

Delete a Tenant

You can delete a specific tenant.

About This Task

Follow this procedure to delete a tenant.

Note
For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

Run the **`efa tenant delete`** command.

Example

- The following example deletes a specified tenant:

```
(efa:root)root@node-2:~# efa tenant delete --name tenant11

Tenant deleted successfully.

--- Time Elapsed: 233.713805ms ---
```

- The following example deletes a specified tenant even when the EPG is associated with tenant:

```
(efa:root)root@node-2:~# efa tenant delete --name tenant11 --force

Tenant delete with force will delete associated Vrfs, EndpointGroups and PortChannels.
Do you want to proceed (Y/N): y

Tenant deleted successfully.

--- Time Elapsed: 1.999257174s ---
```

Configure a Tenant

You can configure tenant in a fabric.

About This Task

Complete the following tasks to configure a tenant in your XCO fabric:

Procedure

1. Create a Private Tenant on page 152
2. Create a Shared Tenant on page 153
3. Scalability on page 154
4. VLAN-based Tenant on page 154
5. Bridge domain-based Tenant on page 154

Create a Private Tenant

You can create a private tenant. Other tenants cannot use the private tenant resource. Default value of a tenant type is private.

About This Task

Follow this procedure to create a private tenant.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To configure a private tenant, run the following command:

```
efa tenant create --name <tenant-name> --type <type-of-tenants> --vlan-range <vlan-range>
--vrf-count <num-of-vrfs> --port <port-list>
```

Example

- The following example creates a specified tenant with type private:

```
(efa:root)root@node-2:~# efa tenant create --name "tenant11" --type private --vlan-
range 2001-2150 --vrf-count 100 --port 10.20.246.5[0/1-10],10.20.246.6[0/1-10].6[
Tenant created successfully.

--- Time Elapsed: 200.022138ms ---
```

- The following example creates a specified tenant with default type:

```
(efa:root)root@node-2:~# efa tenant create --name "tenant12" --vlan-range 2001-2150 --
vrf-count 100 --port 10.20.246.5[0/13],10.20.246.6[0/13]

Tenant created successfully.

--- Time Elapsed: 277.145486ms ---

Show tenant details

(efa:root)root@node-2:~# efa tenant show --detail
================================================
Name            :  tenant11
Type            :  private
Description     :
VLAN Range      :  2001-2150
L2VNI Range     :
L3VNI Range     :
VRF Count       :  100
Enable BD       :  false
Ports           :  10.20.246.5[0/1-10]
                :  10.20.246.6[0/1-10]

================================================
================================================
Name            :  tenant12
Type            :  private
Description     :
VLAN Range      :  2001-2150
L2VNI Range     :
L3VNI Range     :
VRF Count       :  100
Enable BD       :  false
Ports           :  10.20.246.5[0/13]
                :  10.20.246.6[0/13]
```

```
==========================================
==========================================
Name            :  tenant22
Type            :  shared
Description     :
VLAN Range      :  2-21
L2VNI Range     :
L3VNI Range     :
VRF Count       :  100
Enable BD       :  false
Ports           :  10.20.246.5[0/11]
                :  10.20.246.6[0/11]

==========================================
--- Time Elapsed: 72.581191ms ---
```

Create a Shared Tenant

You can create a shared tanant. Other tenants can share the shared tenant, such as VRF and port channel.

About This Task

Follow this procedure to create a shared tenant.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

Run the **`efa tenant create`** command.

```
efa tenant create --name <epg-name> --type shared --port <port-list> --vrf-count <num-
of-vrfs> --l3-vni-range <l3-vni-range> --vlan-range <vlan-range> --l2-vni-range <l2-vni-
range>
```

Example

The following example creates a shared tenant:

```
(efa:root)root@node-2:~# efa tenant create --name "tenant22" --type shared --vlan-range
2-21 --vrf-count 100 --port 10.20.246.5[0/11],10.20.246.6[0/11]46.6[0/11]

Tenant created successfully.

--- Time Elapsed: 223.03097ms ---

Show shared tenant

==========================================================================================
==========================================================================
Name            :  tenant22
Type            :  shared
Description     :
VLAN Range      :  2-21
L2VNI Range     :
L3VNI Range     :
VRF Count       :  100
Enable BD       :  false
Ports           :  10.20.246.5[0/11]
                :  10.20.246.6[0/11]

==========================================================================================
==========================================================================
```

```
--- Time Elapsed: 72.581191ms ---
```

Scalability

The following table provides the scale details for auto and non-auto VNI (Virtual Network Identifier) mapping in a fabric:

Table 15: VNI Scalability

VNI Type	Scale
Non-auto VNI mapping	• The number of VNI (networks) supported per device = 8K [4K VLAN + 4K BD] • The maximum number of VNI (networks) supported in a fabric = [8K * number of devices in the fabric]
Auto VNI mapping	• The number of VNI (networks) supported per device = 8K [4K VLAN + 4K BD] • The number of VNI (networks) supported per fabric = 8K

VLAN-based Tenant

For a VLAN-based tenant, realization of network on the device is done using VLAN and switchport VLANs. Bridge domains are used for EVPN IRB.

The following example creates a VLAN-based tenant with manual VNI mapping:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant11 --l2-vni-range
10002-14190
--l3-vni-range 14191-14200 --vlan-range 2-4090 --vrf-count 10 --port
10.20.216.15[0/11-20],10.20.216.16[0/11-20]
--description Subscriber1

Tenant created successfully.

--- Time Elapsed: 455.141597ms ---
```

Bridge domain-based Tenant

For a BD-based tenant, realization of network on the device is done using BD and BD-LIF. BD is used for EVPN IRB.

The following example creates a BD-based tenant:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant21 --l2-vni-range

30002-34190 --l3-vni-range 34191-34200 --vlan-range 2-4090 --vrf-count 10 --enable-bd --
port 10.20.216.15[0/21-28],10.20.216.16[0/21-28]

Tenant created successfully.

--- Time Elapsed: 501.176996ms ---
```

Provision a Port Channel

You can configure port channels in a fabric.

About This Task

Complete the following tasks to configure a port channel in your XCO fabric:

Procedure

1. Create a Port Channel on page 155
2. Update a Port Channel on page 157
3. Delete a Port Channel on page 159
4. Show a Port Channel on page 160
5. Configure a Port Channel on page 162

Create a Port Channel

You can create a port channel for a tenant. A port channel, also known as a Link Aggregation Group (LAG) is a communication link between devices. You can specify speed, LACP negotiation, port, port channel number, LACP timeout, and the number of links that are required to be up.

About This Task

Follow this procedure to create a port channel.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

Run the **`efa tenant po create`** command to create a port channel.

```
efa tenant po create --name <po-name> --tenant <tenant-name> --description <po-
description> --mtu <1500-9216> --speed <100Mbps|1Gbps|10Gbps|25Gbps|40Gbps|100Gbps> --
negotiation <active|passive|static> --port <list-of-po-members> --min-link-count <min-
link-count> --number <po-number> --lacp-timeout <short|long>
```

Example

- The following example creates a dual-homed PO:

```
efa tenant po create --tenant ten1 --name ten1po1 --port
10.20.246.5[0/1-2],10.20.246.6[0/1-2] --speed 100Gbps --negotiation active --lacp-
timeout short --min-link-count 2 --mtu 9000 --description "ten1po1 of ten1"

PortChannel created successfully.

--- Time Elapsed: 3.291185439s ---


(efa:root)root@node-2:~# efa tenant po show --name ten1po1 --tenant ten1 --detaill
================================================
Name            : ten1po1
Tenant          : ten1
ID              : 7
Description     : ten1po1 of ten1
Speed           : 100Gbps
MTU             : 9000
```

```
Negotiation     : active
Min Link Count  : 2
Lacp Timeout    : short
Ports           : 10.20.246.6[0/1-2]
                : 10.20.246.5[0/1-2]
State           : po-created
Dev State       : provisioned
App State       : cfg-in-sync

================================================

--- Time Elapsed: 99.584783ms ---
```

- The following example shows the SLX configuration on device 1:

```
SSH session to admin@10.20.246.5

SLX# sh run int po
interface Port-channel 1
 speed 100000
 minimum-links 2
 mtu 9000
 description ten1po1 of ten1
 no shutdown
!

SLX# show running-config interface Ethernet 0/1-2
interface Ethernet 0/1
 description Port-channel ten1po1 Member interface
 channel-group 1 mode active type standard
 lacp timeout short
 no shutdown
!
interface Ethernet 0/2
 description Port-channel ten1po1 Member interface
 channel-group 1 mode active type standard
 lacp timeout short
 no shutdown
!
```

- The following example creates a single-homed PO:

```
efa tenant po create --tenant "ten1" --name "ten1po2" --port 10.20.246.5[0/9] --speed
10Gbps --negotiation static --min-link-count 1 --description po2

PortChannel created successfully.

--- Time Elapsed: 3.065422112s ---

=====================================================
Name            : ten1po2
Tenant          : ten1
ID              : 6
Description     : po2
Speed           : 10Gbps
MTU             :
Negotiation     : static
Min Link Count  : 1
Lacp Timeout    :
Ports           : 10.20.246.5[0/9]
State           : po-created
Dev State       : provisioned
App State       : cfg-in-sync

=====================================================
```

```
--- Time Elapsed: 171.910633ms ---
```

- The following example shows an SLX configuration on devices:

```
SLX# sh run int po
interface Port-channel 1
 description po2
 no shutdown
!
```

Update a Port Channel

You can update an existing port channel for a tenant. You can update the port-add, port-delete, lacp-timeout, description, min-link-count, mtu-add, and the mtu-delete operations.

About This Task

Follow this procedure to update a port channel.

Procedure

Run the **`efa tenant po update`** command.

```
efa tenant po update --name <po-name> --tenant <tenant-name> --operation <port-add|port-
delete|lacp-timeout|description|min-linkcount> --port <list-of-po-members> --lacp-timeout
string <short|long> --minlink-count <min-link-count> --description <po-description>
```

Example

```
(efa:root)root@node-2:~# efa tenant po create --name ten1po1 --tenant ten1 --port
        10.20.246.5[0/1],10.20.246.6[0/1] --speed 10Gbps --negotiation active --
description
        tenant1po1

PortChannel created successfully.
--- Time Elapsed: 8.900145166s ---


(efa:root)root@node-2:~# efa tenant po show --name ten1po1 --tenant ten1 --detaill
================================================================================
Name             : ten1po1
Tenant           : ten1
ID               : 1
Description      : tenant1po1
Speed            : 10Gbps
MTU              :
Negotiation      : active
Min Link Count   : 1
Lacp Timeout     : long
Ports            : 10.20.246.6[0/1]
                 : 10.20.246.5[0/1]
State            : po-created
Dev State        : provisioned
App State        : cfg-in-sync
================================================================================
--- Time Elapsed: 43.521043ms ---
```

1. Update MTU for Port Channel

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
mtu-add --mtu 5000

PortChannel: ten1po1 updated successfully.
```

```
--- Time Elapsed: 1.432588278s ---
```

2. Update Lacp-timeout for Port Channel

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
lacp-timeout --lacp-timeout short

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 268.24828ms ---
```

3. Update Port for Port Channel

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
port-add --port 10.20.246.5[0/2],10.20.246.6[0/2]

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 1.765536812s ---
```

4. Update mini-link-count for Port Channel

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
min-link-count --min-link-count 2

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 1.396798321s ---
```

5. Show Port Channel details

```
(efa:root)root@node-2:~# efa tenant po show --name ten1po1 --tenant ten1 --detail
================================================================================
====
Name             : ten1po1
Tenant           : ten1
ID               : 1
Description      : tenant1po1
Speed            : 10Gbps
MTU              : 5000
Negotiation      : active
Min Link Count   : 2
Lacp Timeout     : short
Ports            : 10.20.246.6[0/1-2]
                 : 10.20.246.5[0/1-2]
State            : po-created
Dev State        : provisioned
App State        : cfg-in-sync

================================================================================
====

--- Time Elapsed: 68.366068ms ---
```

6. Update delete MTU for Port Channel

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
mtu-delete

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 1.389710725s ---
```

7. Update mini-link-count for Port Channel

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
min-link-count --min-link-count 1

PortChannel: ten1po1 updated successfully.
```

```
--- Time Elapsed: 1.371611014s ---
```

8. **Update delete Port from Port Channel**

```
(efa:root)root@node-2:~# efa tenant po update --name ten1po1 --tenant ten1 --operation
port-delete --port 10.20.246.5[0/1],10.20.246.6[0/1]

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 1.611562693s ---
```

9. **Show Port Channel details**

```
(efa:root)root@node-2:~# efa tenant po show --name ten1po1 --tenant ten1 --detail
================================================================================================
====
Name             : ten1po1
Tenant           : ten1
ID               : 1
Description      : tenant1po1
Speed            : 10Gbps
MTU              :
Negotiation      : active
Min Link Count   : 1
Lacp Timeout     : short
Ports            : 10.20.246.6[0/2]
                 : 10.20.246.5[0/2]
State            : po-created
Dev State        : provisioned
App State        : cfg-in-sync

================================================================================================
====

--- Time Elapsed: 38.90523ms ---
```

Delete a Port Channel

You can delete a port channel.

About This Task

Follow this procedure to delete a port channel.

Note
For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

Run the **`efa tenant po delete`** command.

```
efa tenant po delete [ --name |--force |--tenant |--help]
```

Example

1. The following example deletes the specified POs:

```
efa tenant po delete --name po1,po2 --tenant tenant11
PortChannel: po1 deleted successfully.
PortChannel: po2 deleted successfully.
--- Time Elapsed: 1.133774283s ---
```

2. The following example deletes a PO even when it is associated with an EPGs:

```
efa tenant po delete --name po1 --tenant tenant11 --force
PortChannel Delete with force will update associated EndpointGroups and Networks and
deletes them if there are no other ports associated to them (N/Y): y
PortChannel: po1 deleted successfully.
--- Time Elapsed: 1.890092303s ---
```

Show a Port Channel

You can view a brief or detailed output of the port channel of all tenants, a given tenant, or a given port channel.

About This Task

Follow this procedure to view a port channel configuration details.

Procedure

Run the **`efa tenant po show`** command.

```
efa tenant po show [--name po-name|--tenant tenant-name | --detail |-- help]
```

Example

1. The following example shows brief output of all port channels:

```
$ efa tenant po show
+------+----------+----+--------------------+---------+-------------+--------------+
| Name | Tenant   | ID | Description        | Speed   | Negotiation | MinLinkCount |
+------+----------+----+--------------------+---------+-------------+--------------+
| po1  | tenant11 | 1  | EFA Port-channel po1 | 100Gbps | active      | 2            |
|      |          |    |                    |         |             |              |
+------+----------+----+--------------------+---------+-------------+--------------+
| po2  | tenant11 | 2  | po2                | 10Gbps  | static      | 1            |
|      |          |    |                    |         |             |              |
+------+----------+----+--------------------+---------+-------------+--------------+
| po11 | tenant21 | 3  |EFA Port-channel po11 | 25Gbps  | active      | 1            |
|      |          |    |                    |         |             |              |
+------+----------+----+--------------------+---------+-------------+--------------+
+----------------------+-------------+------------+-------------+-------------+
| Ports                | LacpTimeout | State      | Dev-State   | App-State   |
+----------------------+-------------+------------+-------------+-------------+
| 10.20.216.15[0/12-13] | short       | po-created | provisioned | cfg-in-sync |
| 10.20.216.16[0/12-13] |             |            |             |             |
+----------------------+-------------+------------+-------------+-------------+
| 10.20.216.15[0/15]    |             | po-created | provisioned | cfg-in-sync |
| 10.20.216.16[0/15]    |             |            |             |             |
+----------------------+-------------+------------+-------------+-------------+
| 10.20.216.15[0/22]    | short       | po-created | provisioned | cfg-in-sync |
| 10.20.216.16[0/22]    |             |            |             |             |
+----------------------+-------------+------------+-------------+-------------+
PortChannel Details
--- Time Elapsed: 832.496716ms ---
```

2. The following example shows detailed output of all port channels:

```
$ efa tenant po show --detail
================================================================================================
Name : po1
Tenant : tenant11
ID : 1
Description : EFA Port-channel po1
Speed : 100Gbps
Negotiation : active
```

```
Min Link Count : 2
Lacp Timeout : short
Ports : 10.20.216.15[0/12-13]
 : 10.20.216.16[0/12-13]
State : po-created
Dev State : provisioned
App State : cfg-in-sync
================================================================================
================================================================================
Name : po2
Tenant : tenant11
ID : 2
Description : EFA Port-channel po3
Speed : 10Gbps
Negotiation : static
Min Link Count : 1
Lacp Timeout :
Ports : 10.20.216.15[0/15]
 : 10.20.216.16[0/15]
State : po-created
Dev State : provisioned
App State : cfg-in-sync
================================================================================
================================================================================
Name : po11
Tenant : tenant21
ID : 3
Description : EFA Port-channel po11
Speed : 25Gbps
Negotiation : active
Min Link Count : 1
Lacp Timeout : short
Ports : 10.20.216.15[0/22]
 : 10.20.216.16[0/22]
State : po-created
Dev State : provisioned
App State : cfg-in-sync
================================================================================
--- Time Elapsed: 506.117046ms ---
```

3. The following example shows brief output of a specific port channels:

```
$ efa tenant po show --tenant tenant11 --name po1
+------+----------+----+---------+-------------+----------+---------+
| Name | Tenant   | ID | Speed   | Negotiation | Min Link | Lacp    |
|      |          |    |         |             | Count    | Timeout |
+------+----------+----+---------+-------------+----------+---------+
| po1  | tenant11 | 1  | 100Gbps | active      | 1        | short   |
|      |          |    |         |             | |        |
+------+----------+----+---------+-------------+----------+---------+
+---------+--------------------+------------+-------------+-------------+
| Lacp | Ports | State | Dev State | App State || Timeout | | | | |
+---------+--------------------+------------+-------------+-------------+
| short | 10.20.216.15[0/12] | po-created | provisioned | cfg-in-sync |
| | 10.20.216.16[0/12] | | | |
+---------+--------------------+------------+-------------+-------------+
PortChannel Details
--- Time Elapsed: 150.30883ms ---
```

4. The following example shows detailed output of all port channels of a tenant:

```
$ efa tenant po show --tenant tenant21 --detail
================================================================================
Name : po11
Tenant : tenant21
ID : 3
Description : EFA Port-channel po11
```

```
Speed : 25Gbps
Negotiation : active
Min Link Count : 1
Lacp Timeout : short
Ports : 10.20.216.15[0/22]
 : 10.20.216.16[0/22]
State : po-created
Dev State : provisioned
App State : cfg-in-sync
================================================================================
--- Time Elapsed: 223.892847ms ---
```

Configure a Port Channel

You can configure a port channel for a tenant. Use the **`efa tenant po configure`** command to push or remove the pending port channel configuration. The command pushes the pending configuration for a port channel when the port channel is in `po-delete-pending`, `port-delete-pending` or `po-lacp-timeout-set-pending` state.

About This Task

Follow this procedure to configure a port channel.

Procedure

Run the **`efa tenant po configure`** command.

```
efa tenant po configure [ --name | --tenant | --help ]
```

Example

The following example pushes or removes the pending port channel configuration:

```
$ efa tenant po show --name ten1po1 --tenant ten1 --detail
============================================================
Name                : ten1po1
Tenant              : ten1
ID                  : 1
Description         : tenant1po1
Speed               : 10Gbps
MTU                 :
Negotiation         : active
Min Link Count      : 2
Lacp Timeout        : long
Ports               : 10.20.246.6[0/1-2]
                    : 10.20.246.5[0/1-2]
State               : port-delete-pending
Dev State           : provisioned
App State           : cfg-in-sync

============================================================

--- Time Elapsed: 111.335777ms ---


$ efa tenant po configure --name ten1po1 --tenant ten1

PortChannel: ten1po1 configured successfully.

--- Time Elapsed: 114.816703ms ---
```

```
$ efa tenant po show --name ten1po1 --tenant ten1 --detail
=========================================================
Name              : ten1po1
Tenant            : ten1
ID                : 1
Description       : tenant1po1
Speed        : 10Gbps
MTU               :
Negotiation       : active
Min Link Count    : 2
Lacp Timeout      : long
Ports        : 10.20.246.6[0/1-2]
                   : 10.20.246.5[0/1-2]
State        : po-created
Dev State         : provisioned
App State         : cfg-in-sync

=========================================================

--- Time Elapsed: 120.391994ms ---
```

Configure Description on Port Channel

You can configure description for each XCO port channel when you create or update a port channel. The default value of a port channel "description" is **"EFA Port-channel <efa-po-name>"**.

About This Task

Follow this procedure to configure description on a port channel.

Procedure

1. Run the following command to configure description when you create a port channel:

```
efa tenant po create --name <po-name> --tenant <tenant-name> --description <po-
description>
        --speed <100Mbps|1Gbps|10Gbps|25Gbps|40Gbps|100Gbps> --negotiation <active|
passive|static>
   --port <list-of-po-members> --min-link-count <min-link-count> --number <po-number>
--lacp-timeout <short|long>
```

2. Run the following command to configure description when you update a port channel:

```
efa tenant po update --name <po-name> --tenant <tenant-name>
    --operation <port-add|port-delete|lacp-timeout|description|min-link-count>
    --port <list-of-po-members> --lacp-timeout string <short|long> --min-link-count
<min-link-count>
    --description <po-description>
```

   The following example shows configuration of description attribute when you create or update a port channel:

```
efa tenant po create --name ten1po1 --tenant ten1 --port
10.20.246.15[0/1],10.20.246.16[0/1] --speed 10Gbps --negotiation active --description
tenant1po1

efa tenant po create --name ten1po2 --tenant ten1 --port
10.20.246.15[0/2],10.20.246.16[0/2] --speed 10Gbps --negotiation active
```

```
efa tenant po update --name ten1po1 --tenant ten1 --operation description --
description tenat1po1changed
```

efa tenant po show --name ten1po1 --tenant ten1 --detail

```
Name              : ten1po1
Tenant            : ten1
ID                : 1
Description       : tenat1po1changed
Speed             : 10Gbps
Negotiation       : active
Min Link Count    : 1
Lacp Timeout      : long
Ports             : 10.20.246.15[0/1]
                  : 10.20.246.16[0/1]
State             : po-created
Dev State         : provisioned
App State         : cfg-in-sync
```

efa tenant po show --name ten1po2 --tenant ten1 --detail

```
Name              : ten1po2
Tenant            : ten1
ID                : 2
Description       : EFA Port-channel
ten1po2
Speed             : 10Gbps
Negotiation       : active
Min Link Count    : 1
Lacp Timeout      : long
Ports             : 10.20.246.15[0/2]
                  : 10.20.246.16[0/2]
State             : po-created
Dev State         : provisioned
App State         : cfg-in-sync
```

Configure Minimum Link Count on Port Channel

You can configure minimum number of link on a port channel. You can provide an optional "min-link-count" for each XCO port channel when you create or update a port channel.

About This Task

Follow this procedure to configure minimum link count on a port channel.

Default value of minimum link count (min-link-count) for a port-channel is 1, which is equal to the SLX default value. When you update the min-link-count attribute, XCO validates the port count on port channel member and minimum link count on each device.

Note

During upgrade from EFA 2.5.5 to the above versions of EFA, the min-link-count for the port-channels is set to the default value 1.

- **Empty port channel:** EFA 2.5.5 and above does not support empty port channel. Therefore, during upgrade from EFA 2.5.5 to the above versions of EFA, all the empty port channels are marked with “delete-pending” state.
- **Non-empty port channel:** During upgrade from EFA 2.5.5 to the above versions of EFA, all the non-empty port channels get configured with the default value (1) of `min-link-count`, and displayed in the **`efa tenant po show`** command output.
- Single Homed to Dual Homed port channel conversion is not allowed in EFA 2.5.5 and above.

Procedure

1. To configure minimum link count when you create a port channel, run the following command:

```
efa tenant po create --name <po-name> --tenant <tenant-name> --description <po-
description>
   --speed <100Mbps|1Gbps|10Gbps|25Gbps|40Gbps|100Gbps> --negotiation <active|passive|
static>
   --port <list-of-po-members> --min-link-count <min-link-count>
   --number <po-number> --lacp-timeout <short|long>
```

2. To configure minimum link count when you update a port channel, run the following command:

```
efa tenant po update --name <po-name> --tenant <tenant-name>
    --operation <port-add|port-delete|lacp-timeout|description|min-link-count>
    --port <list-of-po-members> --lacp-timeout string <short|long> --min-link-count
<min-link-count>
    --description <po-description>
```

   The following example configures minimum link count during port channel create and update operations:

```
efa tenant po create --name ten1po1 --tenant ten1 --port
10.20.246.15[0/1-2],10.20.246.16[0/1-2]
--speed 10Gbps --negotiation active --description tenant1po1 --min-link-count 2

efa tenant po create --name ten1po2 --tenant ten1 --port
10.20.246.15[0/3],10.20.246.16[0/3]
--speed 10Gbps --negotiation active

efa tenant po update --name ten1po1 --tenant ten1 --operation port-delete --port
10.20.246.15[0/1],10.20.246.16[0/1] --min-link-count 1

efa tenant po update --name ten1po1 --tenant ten1 --operation port-add --port
10.20.246.15[0/1],10.20.246.16[0/1] --min-link-count 2

efa tenant po update --name ten1po1 --tenant ten1 --operation min-link-count --min-
link-count 1

efa tenant po update --name ten1po1 --tenant ten1 --operation min-link-count --min-
link-count 2

efa tenant po show
+-------+------+---+------+-----------+--------+-------+--------------------
+-----------+-----------+-----------+
| Name  |Tenant|ID |Speed |Negotiation|Min Link| Lacp  |
Ports        |  State    | Dev State | App State |
|       |      |   |      |           | Count  |Timeout|
|           |           |           |
+-------+------+---+------+-----------+--------+-------+--------------------
+-----------+-----------+-----------+
|ten1po1| ten1 | 1 |10Gbps|  active   |   2
| long  |10.20.246.15[0/1-2] |po-created |provisioned|cfg-in-sync|
|       |      |   |      |           |        |
|10.20.246.16[0/1-2] |           |           |           |
+-------+------+---+------+-----------+--------+-------+--------------------
+-----------+-----------+-----------+
|ten1po2| ten1 | 2 |10Gbps|  active   |   1
| long  | 10.20.246.15[0/3]  |po-created |provisioned|cfg-in-sync|
|       |      |   |      |           |        |
| 10.20.246.16[0/3]  |           |           |           |
+-------+------+---+------+-----------+--------+-------+--------------------
+-----------+-----------+-----------+
```

Configure MTU on Port Channel

You can provide an MTU value on each port channel when you create or update a port channel.

About This Task

Follow this procedure to configure a port channel.

If you do not provide an MTU value, depending on the global MTU configuration, SLX determines a default value of MTU on port channel. If the global MTU is configured, then the MTU value of a port channel inherits the global MTU value. If you have not configured a global MTU, SLX determines a default value of MTU on port channel..

Note
When you configure an ethernet port as a port channel member with an MTU value, the create or update operation of port channel with this ethernet port fails with an appropriate error. Remove the MTU configuration from the ethernet port and then re-attempt the port channel create or update operation.

Procedure

1. To configure an MTU when you create a port channel, run the following command:

```
efa tenant po update --name <po-name> --tenant <tenant-name>
   --operation <port-add|port-delete|lacp-timeout|description|min-link-count|mtu-add|
mtu-delete>
   --port <list-of-po-members> --lacp-timeout string <short|long> --min-link-count
<min-link-count>
   --description <po-description> --mtu <1500-9216>
```

2. Verify the switch configuration on SLX devices.

```
Rack1-Device1# show run interface
Port-channel
interface Port-channel 1
 mtu 9000
 no shutdown
!
interface Port-channel 2
 mtu 7000
 no shutdown
!
Rack1-Device1#
```

```
Rack1-Device2# show run interface
Port-channel
interface Port-channel 1
 mtu 9000
 no shutdown
!
interface Port-channel 2
 mtu 7000
 no shutdown
!
Rack1-Device2#
```

Example

```
efa tenant po create --name ten1po1 --tenant ten1 --port
10.20.246.15[0/1],10.20.246.16[0/1] --speed 10Gbps --negotiation active --mtu 9000

efa tenant po create --name ten1po2 --tenant ten1 --port
10.20.246.15[0/2],10.20.246.16[0/2] --speed 10Gbps --negotiation active

efa tenant po update --name ten1po2 --tenant ten1 --operation mtu-add --mtu 5000

efa tenant po update --name ten1po2 --tenant ten1 --operation mtu-delete
```

```
efa tenant po update --name ten1po2 --tenant ten1 --operation mtu-add --mtu 7000
```

```
efa tenant po show --name ten1po1
--tenant ten1 –detail
==================================
Name             : ten1po1
Tenant           : ten1
ID               : 1
MTU              : 9000
Speed            : 10Gbps
Negotiation      : active
Min Link Count   : 1
Lacp Timeout     : long
Ports            :
10.20.246.15[0/1]
                 : 10.20.246.16[0/1]
State            : po-created
Dev State        : provisioned
App State        : cfg-in-sync
==================================
```

```
efa tenant po show --name ten1po2
--tenant ten1 –detail
==================================
Name             : ten1po2
Tenant           : ten1
ID               : 2
MTU              : 7000
Speed            : 10Gbps
Negotiation      : active
Min Link Count   : 1
Lacp Timeout     : long
Ports            : 10.20.246.15[0/2]
                 : 10.20.246.16[0/2]
State            : po-created
Dev State        : provisioned
App State        : cfg-in-sync
```

SLX configuration example

1. The following is an example configuration on SLX device after creating ten1po1 and ten1po2:

```
SLX# sh run int po
interface Port-channel 1
 mtu 9000
 description EFA Port-channel ten1po1
 no shutdown
!
interface Port-channel 2
 description EFA Port-channel ten1po2
 no shutdown
!
```

2. The following is an example configuration on SLX device after updating ten1po2 to 5000:

```
SLX# sh run int po
interface Port-channel 1
 mtu 9000
 description EFA Port-channel ten1po1
 no shutdown
!
interface Port-channel 2
 mtu 5000
 description EFA Port-channel ten1po2
 no shutdown
!
```

3. The following is an example configuration on SLX device after updating ten1po2 delete MTU:

```
SLX# sh run int po
interface Port-channel 1
 mtu 9000
 description EFA Port-channel ten1po1
 no shutdown
!
interface Port-channel 2
 description EFA Port-channel ten1po2
```

```
 no shutdown
!
```

4. The following is an example configuration on SLX device after updating ten1po2 to 7000:

```
SLX# sh run int po
interface Port-channel 1
 mtu 9000
 description EFA Port-channel ten1po1
 no shutdown
!
interface Port-channel 2
 mtu 7000
 description EFA Port-channel ten1po2
 no shutdown
!
```

Configure LACP-timeout on Port Channel

You can provide a `lacp-timeout` value on each XCO port channel configured on the SLX port channel.

About This Task

Follow this procedure to configure a LACP timeout when you create or update a port channel.

If you do not provide a `lacp-timeout` value, the default value of the port channel `lacp-timeout` will be `long`.

Note
When you configure an ethernet port as a port channel member with the `lacp-timeout` value, set the negotiation to active or passive.

Procedure

1. To configure lacp-timeout when you create a port channel, run the following command:

```
efa tenant po create --name <po-name> --tenant <tenant-name> --description
<podescription> --speed <100Mbps|1Gbps|10Gbps|25Gbps|40Gbps|100Gbps> --negotiation
<active|passive|static> --port <list-of-po-members> --min-link-count <min-link-count>
--number <po-number> --lacp-timeout <short|long>
```

2. To configure lacp-timeout when you update a port channel, run the following command:

```
efa tenant po update --name <po-name> --tenant <tenant-name> --operation <port-
add|port-delete|lacp-timeout|description|min-linkcount> --port <list-of-po-members> --
lacp-timeout string <short|long> --minlink-count <min-link-count> --description <po-
description>
```

Example

The following is an example configuration of LACP timeout on a port channel:

```
efa tenant po create --name ten1po1 --tenant ten1 --port
10.20.246.5[0/1],10.20.246.6[0/1] --speed 10Gbps --negotiation active --lacp-timeout long

PortChannel created successfully.

--- Time Elapsed: 10.513257386s ---
```

```
(efa:root)root@node-2:~# efa tenant po show --name ten1po1 --tenant ten1 --detail
================================================================================
Name            : ten1po1
Tenant          : ten1
ID              : 7
Description     : EFA Port-channel ten1po1
Speed           : 10Gbps
MTU             :
Negotiation     : active
Min Link Count  : 1
Lacp Timeout    : short
Ports           : 10.20.246.6[0/1]
                : 10.20.246.5[0/1]
State           : po-created
Dev State       : provisioned
App State       : cfg-in-sync

================================================================================

--- Time Elapsed: 57.382422ms ---

efa tenant po update --name ten1po1 --tenant ten1 --operation lacp-timeout --lacp-timeout
short

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 1.472657838s ---

efa tenant po show --name ten1po1 --tenant ten1 --detail
================================================================================
Name            : ten1po1
Tenant          : ten1
ID              : 1
Description     : EFA Port-channel ten1po1
Speed           : 10Gbps
MTU             :
Negotiation     : active
Min Link Count  : 1
Lacp Timeout    : short
Ports           : 10.20.246.6[0/1]
                : 10.20.246.5[0/1]
State           : po-created
Dev State       : provisioned
App State       : cfg-in-sync

================================================================================

--- Time Elapsed: 54.009354ms ---
```

The following is an example configuration of LACP timeout on SLX device:

After create ten1po1	After update lacp-timeout to short
SLX# sh run int po interface Port-channel 1 description EFA Port-channel ten1po1 no shutdown ! SLX# sh run int eth 0/1 interface Ethernet 0/1 description Port-channel ten1po1 Member interface channel-group 1 mode active type standard lacp timeout long no shutdown !	SLX# sh run int po interface Port-channel 1 description EFA Port-channel ten1po1 no shutdown ! SLX# sh run int eth 0/1 interface Ethernet 0/1 description Port-channel ten1po1 Member interface channel-group 1 mode active type standard lacp timeout short no shutdown !

Configure Port on a Port Channel

You can provide a port number for each XCO port channel configured on the SLX port channel.

About This Task

Follow this procedure to configure a port when you create or update a port channel..

Procedure

1. To configure a port when you create a port channel, run the following command:

```
efa tenant po create --name <po-name> --tenant <tenant-name> --description
<podescription> --speed <100Mbps|1Gbps|10Gbps|25Gbps|40Gbps|100Gbps> --negotiation
<active|passive|static> --port <list-of-po-members> --min-link-count <min-link-count>
--number <po-number> --lacp-timeout <short|long>
```

2. To configure lacp-timeout when you update a port channel, run the following command:

```
efa tenant po update --name <po-name> --tenant <tenant-name> --operation <port-
add|port-delete|lacp-timeout|description|min-linkcount> --port <list-of-po-members> --
lacp-timeout string <short|long> --minlink-count <min-link-count> --description <po-
description>
```

Example

The following is an example configuration of a port channel:

```
efa tenant po create --name ten1po1 --tenant ten1 --port
10.20.246.5[0/1],10.20.246.6[0/1] --speed 10Gbps --negotiation active
PortChannel created successfully.
--- Time Elapsed: 1.964684168s ---


efa tenant po show --name ten1po1 --tenant ten1 --detaill
================================================================================
Name            : ten1po1
Tenant          : ten1
ID              : 7
Description     : EFA Port-channel ten1po1
Speed           : 10Gbps
MTU             :
```

```
Negotiation     : active
Min Link Count  : 1
Lacp Timeout    : long
Ports           : 10.20.246.6[0/1]
                : 10.20.246.5[0/1]
State           : po-created
Dev State       : provisioned
App State       : cfg-in-sync

================================================================================

--- Time Elapsed: 56.120925ms ---

efa tenant po update --name ten1po1 --tenant ten1 --operation port-add --port
10.20.246.5[0/2],10.20.246.6[0/2]

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 3.201643775s ---

efa tenant po show --name ten1po1 --tenant ten1 --detail
================================================================================
Name            : ten1po1
Tenant          : ten1
ID              : 7
Description     : EFA Port-channel ten1po1
Speed           : 10Gbps
MTU             :
Negotiation     : active
Min Link Count  : 1
Lacp Timeout    : long
Ports           : 10.20.246.6[0/1-2]
                : 10.20.246.5[0/1-2]
State           : po-created
Dev State       : provisioned
App State       : cfg-in-sync

================================================================================

--- Time Elapsed: 64.672251ms ---


efa tenant po update --name ten1po1 --tenant ten1 --operation port-delete --port
10.20.246.5[0/1],10.20.246.6[0/1]

PortChannel: ten1po1 updated successfully.

--- Time Elapsed: 1.71277107s ---

efa tenant po show --name ten1po1 --tenant ten1 --detaill
================================================================================
Name            : ten1po1
Tenant          : ten1
ID              : 7
Description     : EFA Port-channel ten1po1
Speed           : 10Gbps
MTU             :
Negotiation     : active
Min Link Count  : 1
Lacp Timeout    : long
Ports           : 10.20.246.6[0/2]
                : 10.20.246.5[0/2]
State           : po-created
Dev State       : provisioned
```

```
App State       : cfg-in-sync

================================================================================

--- Time Elapsed: 64.928169ms ---
```

The following is an example configuration of a port channel on SLX device:

After create ten1po1	After update port-add operation 0/2
SLX# sh run int po interface Port-channel 1 description EFA Port-channel ten1po1 no shutdown ! SLX# sh run int eth 0/1-2 interface Ethernet 0/1 description Port-channel ten1po1 Member interface channel-group 1 mode active type standard lacp timeout long no shutdown ! interface Ethernet 0/2 no shutdown !	SLX# sh run int po interface Port-channel 1 description EFA Port-channel ten1po1 no shutdown ! SLX# sh run int eth 0/1-2 interface Ethernet 0/1 description Port-channel ten1po1 Member interface channel-group 1 mode active type standard lacp timeout long no shutdown ! interface Ethernet 0/2 description Port-channel ten1po1 Member interface channel-group 1 mode active type standard lacp timeout long no shutdown !

Shared and Private Port Channel Configuration

The following is an example configuration of shared and private port channel:

```
efa tenant po create --name sharedPO --tenant sharedTenant
   --port 10.20.246.15[0/31],10.20.246.16[0/31] --speed 10Gbps --negotiation active

efa tenant po create --name ten1po1 --tenant tenant1
   --port 10.20.246.17[0/11],10.20.246.18[0/11] --speed 10Gbps --negotiation active

efa tenant po create --name ten1po2 --tenant tenant1
   --port 10.20.246.25[0/11],10.20.246.26[0/11] --speed 10Gbps --negotiation active
efa tenant po create --name ten2po1 --tenant tenant2
   --port 10.20.246.17[0/21],10.20.246.18[0/21] --speed 10Gbps --negotiation active
efa tenant po create --name ten2po2 --tenant tenant2
   --port 10.20.246.25[0/21],10.20.246.26[0/21] --speed 10Gbps --negotiation active

efa tenant po show
+----------+--------------+----+--------+-------------+----------+---------
+--------------------+------------+-------------+-------------+
|   Name   |    Tenant    | ID | Speed  | Negotiation
| Min Link |  Lacp   |       Ports        |   State    |  Dev State  |  App State  |
|          |              |    |        |            |
Count   | Timeout |                    |            |             |             |
+----------+--------------+----+--------+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| sharedPO | sharedTenant | 1  | 10Gbps |   active
|    1     |  long   | 10.20.246.16[0/31] | po-created | provisioned | cfg-in-sync |
```

```
|          |              |    |        |             |
|          | 10.20.246.15[0/31] |           |             |             |
+----------+--------------+----+--------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
| ten1po1  |   tenant1    | 1  | 10Gbps |   active
|    1     |  long    | 10.20.246.18[0/11] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |             |
|          | 10.20.246.17[0/11] |           |             |             |
+----------+--------------+----+--------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
| ten1po2  |   tenant1    | 1  | 10Gbps |   active
|    1     |  long    | 10.20.246.25[0/11] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |
|          |          | 10.20.246.26[0/11] |            |             |             |
+----------+--------------+----+--------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
| ten2po1  |   tenant2    | 2  | 10Gbps |   active
|    1     |  long    | 10.20.246.18[0/21] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |
|          |          | 10.20.246.17[0/21] |            |             |             |
+----------+--------------+----+--------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
| ten2po2  |   tenant2    | 2  | 10Gbps |   active
|    1     |  long    | 10.20.246.25[0/21] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |
|          |          | 10.20.246.26[0/21] |            |             |             |
+----------+--------------+----+--------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
```

Provision a VRF

You can configure a VRF in a tenant network.

About This Task

Complete the following tasks to configure a tenant VRF in your XCO fabric:

Procedure

1. Create a Tenant VRF on page 173
2. Update a Tenant VRF on page 177
3. Show a Tenant VRF on page 181
4. Delete a Tenant VRF on page 184
5. Shows a Tenant VRF Error on page 185
6. Configure a Tenant VRF on page 185
7. Distributed and Centralized Routing on page 212

Create a Tenant VRF

You can configure virtual routing and forwarding (VRF) for the tenant. You can specify the name of VRF and the associated tenant, target VPN community, Route Target and Route Distinguisher, Local ASN, IPv4 and IPv6 static BFD routes, IPv4 and IPv6 static next hop routes, number of load sharing paths, redistribute type, whether resilient hashing is on SLX devices, and the routing type.

About This Task

Follow this procedure to create a tenant VRF.

Procedure

To configure a tenant VRF, run the following command:

```
efa tenant vrf create [ --name vrf-name | --tenant tenant-name | --rt-type { both |
import | export }| --rt value |--local-asn local-asn | --ipv4-static-route-bfd route |
--ipv6-static-route-bfd route | --ipv4-static-route-next-hop route | --ipv6-static-route-
next-hop route
| --max-path unit |--redistribute { static | connected} | --rh-maxpath { 8 | 16 | 64 }| --
rh-ecmp-enable= {true | false } | --graceful-restart-enable= {true | false } | --routing-
type { distributed |centralized }|--help ]
```

Note

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Example

1. The following example creates a distributed VRF:

```
(efa:root)root@node-2:~# efa tenant vrf create --tenant tenant11 --name blue11 --local-
asn 65001 --rt-type import --rt 100:100 --rt-type export --rt 100:100 --rt-type import
--rt 200:200 --rt-type export --rt 200:200 --rt-type import --rt 300:300 --rt-type
export --rt 400:400 --max-path 50 --redistribute connected --redistribute static --
ipv4-static-route-next-hop 10.20.246.6,192.168.0.0/24,10.10.10.1,5 --ipv4-static-route-
next-hop 10.20.246.5,192.168.10.0/24,10.10.10.5,5 --ipv6-static-route-next-hop
10.20.246.6,2020:20::1/128,3001::2,6 --ipv6-static-route-next-hop
10.20.246.5,2020:30::1/128,3001::3,5 --ipv6-static-route-bfd
10.20.246.6,3001::3,3001::1,100,200,5 --ipv6-static-route-bfd
10.20.246.6,3001::2,3001::1 --ipv6-static-route-bfd
10.20.246.6,3001::4,3001::1,100,300,6 --ipv4-static-route-bfd
10.20.246.5,10.10.10.1,10.10.10.254,200,300,6 --ipv4-static-route-bfd
10.20.246.6,10.10.10.5,10.10.10.252 --rh-ecmp-enable --rh-max-path 16 --graceful-
restart-enable --routing-type distributed

Vrf created successfully.

--- Time Elapsed: 772.62533ms ---

(efa:root)root@node-2:~# efa tenant vrf show --name blue11 --tenant tenant11 --detail
================================================================================
================================================================
Name                     : blue11
Tenant                   : tenant11
Routing Type             : distributed
Centralized Routers      :
Redistribute             : connected,static
Max Path                 : 50
Local Asn                : 65001
L3VNI                    :
EVPN IRB BD              :
EVPN IRB VE              :
BR VNI                   :
BR BD                    :
BR VE                    :
RH Max Path              : 16
Enable RH ECMP           : true
Enable Graceful Restart  : true
Route Target             : import 100:100
                         : export 100:100
                         : import 200:200
```

```
                        : export 200:200
                        : import 300:300
                        : export 400:400
Static Route            : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                        : 10.20.246.6->192.168.0.0/24,10.10.10.1[5]
2020:20::1/128,3001::2[6]
                        : 10.20.246.5->192.168.10.0/24,10.10.10.5[5]
2020:30::1/128,3001::3[5]
Static Route BFD        : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                        : 10.20.246.5->10.10.10.1,10.10.10.254[200,300,6]
                        : 10.20.246.6->10.10.10.5,10.10.10.252
3001::3,3001::1[100,200,5] 3001::2,3001::1 3001::4,3001::1[100,300,6]
VRF Type                :
State                   : vrf-create
Dev State               : not-provisioned
App State               : cfg-ready

================================================================================
================================================================================

--- Time Elapsed: 47.564929ms ---



(efa:root)root@node-2:~# efa tenant epg create --name epg1 --tenant tenant11 --port
10.20.246.5[0/1],10.20.246.6[0/1] --vrf blue11 --switchport-mode trunk --ctag-range
2001 --anycast-ip 2001:10.10.11.1/24

EndpointGroup created successfully.

--- Time Elapsed: 12.400157552s ---
```

2. The following is an example configuration on SLX devices:

On Device1: 10.20.246.5

```
SLX# show running-config vrf
vrf blue11
 rd 172.31.254.211:1
 resilient-hash ecmp enable
 resilient-hash max-path 16
 evpn irb ve 8192
 address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ip route static bfd 10.10.10.1
10.10.10.254 interval 200 min-rx
300 multiplier 6
  ip route 192.168.10.0/24
10.10.10.5 distance 5
 !
 address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ipv6 route 2020:30::1/128
3001::3 distance 5
 !
!
```

On Device2: 10.20.246.6

```
SLX# show running-config vrf
vrf blue11
 rd 172.31.254.152:1
 resilient-hash ecmp enable
 resilient-hash max-path 16
 evpn irb ve 8192
 address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ip route static bfd 10.10.10.5
10.10.10.252
  ip route 192.168.0.0/24
10.10.10.1 distance 5
 !
 address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ipv6 route static bfd 3001::2
3001::1
  ipv6 route static bfd 3001::3
3001::1 interval 100 min-rx 200
multiplier 5
  ipv6 route static bfd 3001::4
3001::1 interval 100 min-rx 300
multiplier 6
  ipv6 route 2020:20::1/128
3001::2 distance 6
 !
!
```

3. The following example creates a centralized VRF:

```
(efa:extreme)extreme@node-1:~$ efa tenant vrf create --name red13 --tenant
tenant21 --max-path 50 --redistribute connected --redistribute static --local-
asn 65002 --ipv4-static-route-next-hop 10.20.216.104,192.168.0.0/24,10.10.10.1,5
--ipv4-static-route-next-hop 10.20.216.104,192.168.10.0/24,10.10.10.5,5
--ipv6-static-route-next-hop 10.20.216.104,2020:20::1/128,3001::2,6
--ipv6-static-route-next-hop 10.20.216.104,2020:30::1/128,3001::3,5 --
ipv6-static-route-bfd 10.20.216.104,3001::3,3001::1,100,200,5 --ipv6-
static-route-bfd 10.20.216.104,3001::2,3001::1 --ipv6-static-route-
bfd 10.20.216.104,3001::4,3001::1,100,300,6 --ipv4-static-route-bfd
10.20.216.104,10.10.10.1,10.10.10.254,200,300,6 --ipv4-static-route-bfd
10.20.216.104,10.10.10.5,10.10.10.252 --rh-max-path 64 --routing-type centralized --
centralized-router 10.20.216.103,10.20.216.104

Vrf created successfully.

--- Time Elapsed: 726.425268ms ---
```

Update a Tenant VRF

You can update an existing VRF for a tenant. You can update operations, such as local-asn-add, local-asn-delete, static-route-bfd-add, static-route-bfd-delete, static-route-add, static-route-delete, max-path-add, max-path-delete, redistribute-add, redistribute-delete, rh-max-path-add, rh-max-path-delete, centralized-router-add, centralized-router-delete, rh-ecmp-update, and graceful-restart-update.

About This Task

Follow this procedure to update a tenant VRF.

Procedure

To update a VRF, run the following command:

```
efa tenant vrf update [--name vrf-name | --tenant tenant-name | -- operation code |--
local-asn local-asn | --ipv4-static-route-bfd route | --ipv6-static-route-bfd route | --
ipv4-static-route-next-hop route | --ipv6-static-route-next-hop route | --max-path unit
|-- redistribute {static | connected} | --rh-max-path {8 | 16 | 64} | -- rh-ecmp-enable=
{true | false} | --graceful-restart-enable= {true | false } | --routing-type {distributed
| centralized }
```

Note

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Example

1. The following example updates a local ASN for VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation local-asn-add --local-asn 75001
WARNING : This operation will result in the reset of the backup routing bgp neighbours
of the VRF. Do you want to proceed [y/n]?
y

Vrf updated successfully.

--- Time Elapsed: 7.09160915s ---
```

2. The following example updates a static route for VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant
tenant11 --name blue11 --operation static-route-add --ipv4-static-route-
next-hop 10.20.246.6,182.20.0.0/24,11.11.11.1,5 --ipv6-static-route-next-hop
10.20.246.5,1010:30::1/128,1001::3,5

Vrf updated successfully.

--- Time Elapsed: 244.090637ms ---
```

3. The following example updates a max-path for VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation max-path-add --max-path 60

Vrf updated successfully.

--- Time Elapsed: 188.793294ms ---
```

4. The following example updates a redistribute attribute for VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation redistribute-add --redistribute connected

Vrf updated successfully.
```

```
--- Time Elapsed: 225.778861ms ---
```

5. The following example updates a rh-max-path for VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation rh-max-path-add --rh-max-path 64

Vrf updated successfully.

--- Time Elapsed: 99.141472ms ---
```

6. The following example updates a rh-ecmp-enable for VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation rh-ecmp-update --rh-ecmp-enable=false

Vrf updated successfully.

--- Time Elapsed: 173.438931ms ---
```

7. The following example shows VRF details:

```
(efa:root)root@node-2:~# efa tenant vrf show --name blue11 --tenant tenant11 --detail
================================================================================
================================================
Name                     : blue11
Tenant                   : tenant11
Routing Type             : distributed
Centralized Routers      :
Redistribute             : connected,static
Max Path                 : 60
Local Asn                : 75001
L3VNI                    :
EVPN IRB BD              :
EVPN IRB VE              :
BR VNI                   :
BR BD                    :
BR VE                    :
RH Max Path              : 64
Enable RH ECMP           : false
Enable Graceful Restart  : false
Route Target             : import 100:100
                         : export 100:100
                         : import 200:200
                         : export 200:200
                         : import 300:300
                         : export 400:400
Static Route             : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                         : 10.20.246.6->192.168.0.0/24,10.10.10.1[5]
2020:20::1/128,3001::2[6] 182.20.0.0/24,11.11.11.1[5]
                         : 10.20.246.5->192.168.10.0/24,10.10.10.5[5]
2020:30::1/128,3001::3[5] 1010:30::1/128,1001::3[5]
Static Route BFD         : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                         : 10.20.246.5->10.10.10.1,10.10.10.254[200,300,6]
                         : 10.20.246.6->10.10.10.5,10.10.10.252
3001::3,3001::1[100,200,5] 3001::2,3001::1 3001::4,3001::1[100,300,6]
VRF Type                 :
State                    : vrf-create
Dev State                : not-provisioned
App State                : cfg-ready

================================================================================
================================================

--- Time Elapsed: 59.390211ms ---
```

8. The following is an example of SLX Configuration:

On Device1: 10.20.246.5

```
SLX# show running-config vrf
vrf blue11
 rd 172.31.254.211:1
 resilient-hash max-path 64
 evpn irb ve 8192
 address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ip route static bfd 10.10.10.1
10.10.10.254 interval 200 min-rx
300 multiplier 6
  ip route 192.168.10.0/24
10.10.10.5 distance 5
 !
 address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ipv6 route 1010:30::1/128
1001::3 distance 5
  ipv6 route 2020:30::1/128
3001::3 distance 5
 !
!
```

On Device2: 10.20.246.6

```
SLX# show running-config vrf
vrf blue11
 rd 172.31.254.152:1
 resilient-hash max-path 64
 evpn irb ve 8192
 address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ip route static bfd 10.10.10.5
10.10.10.252
  ip route 182.20.0.0/24
11.11.11.1 distance 5
  ip route 192.168.0.0/24
10.10.10.1 distance 5
 !
 address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ipv6 route static bfd 3001::2
3001::1
  ipv6 route static bfd 3001::3
3001::1 interval 100 min-rx 200
multiplier 5
  ipv6 route static bfd 3001::4
3001::1 interval 100 min-rx 300
multiplier 6
  ipv6 route 2020:20::1/128
3001::2 distance 6
```

9. The following example deletes a local ASN from VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation local-asn-delete
WARNING : This operation will result in the reset of the backup routing bgp neighbours
of the VRF. Do you want to proceed [y/n]?
y

Vrf updated successfully.

--- Time Elapsed: 1.162426042s ---
```

10. The following example deletes a static route from VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant
tenant11 --name blue11 --operation static-route-delete --ipv4-static-route-
next-hop 10.20.246.6,182.20.0.0/24,11.11.11.1,5 --ipv6-static-route-next-hop
10.20.246.5,1010:30::1/128,1001::3,5

Vrf updated successfully.

--- Time Elapsed: 162.621663ms ---
```

11. The following example deletes a static route BFD from VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant
tenant11 --name blue11 --operation static-route-bfd-delete --ipv6-
```

```
static-route-bfd 10.20.246.6,3001::3,3001::1,100,200,5 --ipv4-static-route-bfd
10.20.246.5,10.10.10.1,10.10.10.254,200,300,6

Vrf updated successfully.

--- Time Elapsed: 168.307373ms ---
```

12. The following example deletes Max Path from VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation max-path-delete

Vrf updated successfully.

--- Time Elapsed: 117.514104ms ---
```

13. The following example deletes Redistribute from VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation redistribute-delete --redistribute connected

Vrf updated successfully.

--- Time Elapsed: 202.742522ms ---
```

14. The following example deletes RH Max Path from VRF:

```
(efa:root)root@node-2:~# efa tenant vrf update --tenant tenant11 --name blue11 --
operation rh-max-path-delete

Vrf updated successfully.

--- Time Elapsed: 138.245305ms ---
```

15. The following example shows VRF details:

```
(efa:root)root@node-2:~# efa tenant vrf show --name blue11 --tenant tenant11 --detaill
================================================================================================
====================
Name                      : blue11
Tenant                    : tenant11
Routing Type              : distributed
Centralized Routers       :
Redistribute              : static
Max Path                  : 0
Local Asn                 :
L3VNI                     :
EVPN IRB BD               :
EVPN IRB VE               :
BR VNI                    :
BR BD                     :
BR VE                     :
RH Max Path               :
Enable RH ECMP            : false
Enable Graceful Restart : false
Route Target              : import 100:100
                          : export 100:100
                          : import 200:200
                          : export 200:200
                          : import 300:300
                          : export 400:400
Static Route              : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                          : 10.20.246.6->192.168.0.0/24,10.10.10.1[5]
2020:20::1/128,3001::2[6]
                          : 10.20.246.5->192.168.10.0/24,10.10.10.5[5]
2020:30::1/128,3001::3[5]
Static Route BFD          : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                          : 10.20.246.6->10.10.10.5,10.10.10.252 3001::2,3001::1
3001::4,3001::1[100,300,6]
```

```
VRF Type                :
State                   : vrf-create
Dev State               : not-provisioned
App State               : cfg-ready

================================================================================
==================

--- Time Elapsed: 75.948924ms ---
```

16. The following is an example of SLX Configuration:

On Device1: 10.20.246.5

```
SLX# show running-config vrf
vrf blue11
 rd 172.31.254.211:1
 evpn irb ve 8192
 address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ip route 192.168.10.0/24
10.10.10.5 distance 5
 !
 address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ipv6 route 2020:30::1/128
3001::3 distance 5
 !
!
```

On Device2: 10.20.246.6

```
SLX# show running-config vrf
vrf blue11
 rd 172.31.254.152:1
 evpn irb ve 8192
 address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ip route static bfd 10.10.10.5
10.10.10.252
  ip route 192.168.0.0/24
10.10.10.1 distance 5
 !
 address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
  ipv6 route static bfd 3001::2
3001::1
  ipv6 route static bfd 3001::4
3001::1 interval 100 min-rx 300
multiplier 6
  ipv6 route 2020:20::1/128
3001::2 distance 6
 !
!
```

Show a Tenant VRF

You can view a brief or detailed output of the VRF of all tenants, a given tenant, or a given VRF.

About This Task

Follow this procedure to view a tenant VRF configuration.

Procedure

Run the **`efa tenant vrf show`** command.

Example

1. The following example shows brief output of all VRFs:

```
(efa:root)root@node-2:~# efa tenant vrf show
+--------+----------+--------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
|  Name  |  Tenant  | Routing Type | Centralized Routers |   Redistribute   | Max Path
| Local Asn | Enable GR |   State    |    Dev State    | App State |
+--------+----------+--------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
| blue11 | tenant11 | distributed  |                    | connected,static |    50
|   65001   |   true    | vrf-create | not-provisioned | cfg-ready |
+--------+----------+--------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
| red11  | tenant11 | distributed  |                    | connected,static |    50
|   5001    |   true    | vrf-create | not-provisioned | cfg-ready |
+--------+----------+--------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
Vrf Details

--- Time Elapsed: 171.561968ms ---
```

2. The following example shows detailed output of all VRF:

```
(efa:root)root@node-2:~# efa tenant vrf show --detail
================================================================================================
================================================
Name                     : blue11
Tenant                   : tenant11
Routing Type             : distributed
Centralized Routers      :
Redistribute             : connected,static
Max Path                 : 50
Local Asn                : 65001
L3VNI                    :
EVPN IRB BD              :
EVPN IRB VE              :
BR VNI                   :
BR BD                    :
BR VE                    :
RH Max Path              : 16
Enable RH ECMP           : true
Enable Graceful Restart : true
Route Target             : import 100:100
                         : export 100:100
                         : import 200:200
                         : export 200:200
                         : import 300:300
                         : export 400:400
Static Route             : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                         : 10.20.246.6->192.168.0.0/24,10.10.10.1[5]
2020:20::1/128,3001::2[6]
                         : 10.20.246.5->192.168.10.0/24,10.10.10.5[5]
2020:30::1/128,3001::3[5]
Static Route BFD         : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                         : 10.20.246.5->10.10.10.1,10.10.10.254[200,300,6]
                         : 10.20.246.6->10.10.10.5,10.10.10.252
3001::3,3001::1[100,200,5] 3001::2,3001::1 3001::4,3001::1[100,300,6]
VRF Type                 : private
State                    : vrf-create
Dev State                : not-provisioned
App State                : cfg-ready
```

```
================================================================================================
================================================
================================================================================================
================================================
Name                     : red11
Tenant                   : tenant11
Routing Type             : distributed
Centralized Routers      :
Redistribute             : connected,static
Max Path                 : 50
Local Asn                : 5001
L3VNI                    :
EVPN IRB BD              :
EVPN IRB VE              :
BR VNI                   :
BR BD                    :
BR VE                    :
RH Max Path              : 16
Enable RH ECMP           : true
Enable Graceful Restart : true
Route Target             : import 500:500
                         : export 500:500
                         : import 600:600
                         : export 600:600
                         : import 700:700
                         : export 800:800
Static Route             : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                         : 10.20.246.6->192.168.0.0/24,10.10.10.1[5]
2020:20::1/128,3001::2[6]
                         : 10.20.246.5->192.168.10.0/24,10.10.10.5[5]
2020:30::1/128,3001::3[5]
Static Route BFD         : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                         : 10.20.246.5->10.10.10.1,10.10.10.254[200,300,6]
                         : 10.20.246.6->10.10.10.5,10.10.10.252
3001::3,3001::1[100,200,5] 3001::2,3001::1 3001::4,3001::1[100,300,6]
VRF Type                 : private
State                    : vrf-create
Dev State                : not-provisioned
App State                : cfg-ready

================================================================================================
================================================

--- Time Elapsed: 192.291858ms ---
```

3. The following example shows brief output of a specific VRF:

```
(efa:root)root@node-2:~# efa tenant vrf show --name blue11 --tenant tenant11
+--------+----------+-------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
|  Name  |  Tenant  | Routing Type | Centralized Routers |   Redistribute   | Max Path
| Local Asn | Enable GR |   State    |    Dev State    | App State |
+--------+----------+-------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
| blue11 | tenant11 | distributed  |                    | connected,static |    50
|   65001   |   true    | vrf-create | not-provisioned | cfg-ready |
+--------+----------+-------------+--------------------+------------------+----------
+-----------+-----------+------------+-----------------+-----------+
Vrf Details

--- Time Elapsed: 59.752192ms ---
```

4. The following example shows detailed output of specific VRF of a tenant:

```
(efa:root)root@node-2:~# efa tenant vrf show --name blue11 --tenant tenant11  --detail
================================================================================================
```

```
================================================
Name                     : blue11
Tenant                   : tenant11
Routing Type             : distributed
Centralized Routers      :
Redistribute             : connected,static
Max Path                 : 50
Local Asn                : 65001
L3VNI                    :
EVPN IRB BD              :
EVPN IRB VE              :
BR VNI                   :
BR BD                    :
BR VE                    :
RH Max Path              : 16
Enable RH ECMP           : true
Enable Graceful Restart  : true
Route Target             : import 100:100
                         : export 100:100
                         : import 200:200
                         : export 200:200
                         : import 300:300
                         : export 400:400
Static Route             : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                         : 10.20.246.6->192.168.0.0/24,10.10.10.1[5]
2020:20::1/128,3001::2[6]
                         : 10.20.246.5->192.168.10.0/24,10.10.10.5[5]
2020:30::1/128,3001::3[5]
Static Route BFD         : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                         : 10.20.246.5->10.10.10.1,10.10.10.254[200,300,6]
                         : 10.20.246.6->10.10.10.5,10.10.10.252
3001::3,3001::1[100,200,5] 3001::2,3001::1 3001::4,3001::1[100,300,6]
VRF Type                 :
State                    : vrf-create
Dev State                : not-provisioned
App State                : cfg-ready

================================================================================================
================================================

--- Time Elapsed: 58.788211ms ---
```

Delete a Tenant VRF

You can delete the VRF for a tenant.

About This Task

Follow this procedure to delete a tenant VRF.

Note

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To delete a tenant VRF, run the following command:

```
efa tenant vrf delete [--name vrf name |--tenant tenant name ]
```

Example

The following example deletes the specified tenant VRF:

```
(efa:root)root@node-2:~# efa tenant vrf delete --name red11 --tenant tenant11

Vrf: red11 deleted successfully.

--- Time Elapsed: 165.470132ms ---
```

Shows a Tenant VRF Error

You can view errors in a configuration of a Tenant VRF.

About This Task

Follow this procedure to view errors in a tenant VRF.

Note
For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To show a tenant VRF error, run the following command:.

```
efa tenant vrf error show [ --name vrf-name | --tenant tenant-name ]
```

Example

The following example shows output of VRF errors for a specific tenant:

```
efa tenant vrf error show --tenant tenant11 --name blue11
============================================================
Name : blue11
Tenant : tenant11
Errors
+--------------+------------------------------------------------------------+
| MgmtIp       | ErrorList                                                  |
+--------------+------------------------------------------------------------+
| 10.20.246.5  | Configure RemoteAsn under Router BGP failed for Vrf :  |
|              | blue11 due to Netconf <x> error                        |
+--------------+------------------------------------------------------------+
| 10.20.246.6  | Configure RemoteAsn under Router BGP failed for Vrf :  |
|              | blue11 due to Netconf <x> error                        |
+--------------+------------------------------------------------------------+
--- Time Elapsed: 195.971ms ---
```

Configure a Tenant VRF

You can configure a tenant VRF.

About This Task

Complete the following tasks to configure a tenant VRF in your XCO fabric:

Procedure

1. Configure Local ASN on Tenant VRF on page 186
2. Enable Graceful Restart on Tenant VRF on page 193
3. Configure MaxPaths on Tenant VRF on page 194

4. Configure Resilient Hashing on Tenant VRF on page 195
5. Configure Redistribute Attribute on Tenant VRF on page 196
6. Configure Advertise Network and Static Network on Tenant VRF on page 197
7. Configure Aggregate Address on Tenant VRF on page 199
8. Configure EVPN IRB VE Cluster Gateway on a Tenant VRF on page 201
9. Route Distinguisher (RD) Allocation Independent of Route Target (RT) on page 206
10. IPv6 Anycast Gateway Support on page 207
11. Configure Static VRF Route on page 207
12. Configure BFD on Static VRF Route on page 209
13. Configure Backup Routing on Tenant VRF on page 209
14. Distributed and Centralized Routing on page 212
15. BFD Timers for Router BGP BFD and Static Route BFD Sessions on page 222

Configure Local ASN on Tenant VRF

About This Task

Follow this procedure to configure local ASN.

Procedure

1. To configure local ASN when you create a tenant VRF, run the following commands:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name> --local-asn <local-as-for-vrf>
```

2. To configure local ASN on an existing tenant VRF, run the following commands:

```
efa tenant vrf update --name <vrf-name> --tenant <tenant-name> --operation <local-asn-add|local-asn-delete> --local-asn <value>
```

Note
Ensure that the local ASN support on IPv6 AF is checked.

Update Local ASN on VRF

EFA 2.5.5 or above supports update of local ASN on a tenant VRF which is already used in an endpoint group.

Backup Routing

- XCO automates the backup routing configuration among the MCT nodes by configuring IPv4 or IPv6 IBGP neighborship between the MCT nodes.
- When the local ASN for a VRF (used in an endpoint group) is updated using the local-asn-add operation, the remote-asn of the backup routing IPv4 or IPv6 IBGP neighbors also gets updated.
- When the remote ASN of an existing backup routing BGP neighbor is updated, the corresponding BGP session is reset using the `clear ip bgp neighbor <neighbor-ip> vrf <tenant-vrf-name>` command, which lead to traffic disruption till the session is up.
- When the local ASN for a VRF (used in an endpoint group) is deleted using local-asn-delete operation, the remote-asn of the backup routing IPv4 or IPv6 IBGP neighbors

also gets updated to the local-asn configured at the global router bgp level followed by the backup routing bgp session reset.

Ensure that the BGP neighbors update their remote ASN based on the updated local ASN.

Configure Local ASN during VRF Create

You can configure a local ASN when you create a VRF.

About This Task

Follow this procedure to configure a local ASN.

Procedure

1. To configure local ASN when you create a VRF, run the following command:

```
efa fabric show
Fabric Name: default, Fabric Description: Default Fabric, Fabric Stage: 3, Fabric
Type: clos, Fabric Status: created
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE | DEVICE STATE | APP STATE | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+

Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: settings-updated

Updated Fabric Settings: BGP-LL

+-------------+-----+-----------+-------+------+---------------------+-------------
+-------------------+---------------------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME | ASN   | ROLE | DEVICE STATE        | APP STATE   |
CONFIG GEN REASON | PENDING CONFIGS           | VTLB ID | LB ID |
+-------------+-----+-----------+-------+------+---------------------+-------------
+-------------------+---------------------------+---------+-------+
| 10.20.246.1 |     | SLX-1     | 64512 | Spine| provisioned         | cfg in-sync |
NA                | NA                        | NA      | 1     |
| 10.20.246.7 |     | SLX       | 65000 | Leaf | provisioning failed | cfg ready   |
IA,IU,MD,DA       | SYSP-C,MCT-C,MCT-PA,BGP-C,| 2       | 1     |
|             |     |           |       |      |                     |
|                   | INTIP-C,EVPN-C,O-C        |         |       |
| 10.20.246.8 |     | slx-8     | 65000 | Leaf | provisioned         | cfg in-sync |
NA                | NA                        | 2       | 1     |
+-------------+-----+-----------+-------+------+---------------------+-------------
+-------------------+---------------------------+---------+-------+
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password , BFD-RX - Bfd Rx
Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd multiplier,
BFD-ENABLE - Enable Bfd, BGP-MULTIHOP - BGP ebgp multihop, P2PLR - Point-to-Point Link
Range, MCTLR - MCT Link Range, LOIP - Loopback IP Range


CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
```

```
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface
IP, BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete


efa fabric setting show --name fabric1 --advanced | grep -i "backup routing"
| Backup Routing Enable          | Yes                    |
| Backup Routing IPv4 Range      | 10.40.40.0/24          |
| Backup Routing IPv6 Range      | fd40:4040:4040:1::/120 |

efa tenant show
+------+---------+------------+-------------+-------------+-----------+-----------
+----------------------+
| Name |  Type   | VLAN Range | L2VNI Range | L3VNI Range | VRF Count | Enable BD
|        Ports         |
+------+---------+------------+-------------+-------------+-----------+-----------
+----------------------+
| ten1 | private |   11-20    |             |             |    10     |   false   |
10.20.246.16[0/1-10] |
|      |         |            |             |             |           |           |
10.20.246.15[0/1-10] |
+------+---------+------------+-------------+-------------+-----------+-----------
+----------------------+

efa tenant vrf show
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+--------------+
|   Name   | Tenant | Routing Type | Centralized Routers | Redistribute | Max Path |
Local Asn | Enable GR |       State        |  Dev State  |  App State  |
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+--------------+
| ten1vrf1 |  ten1  | distributed  |                     |  connected   |    8     |
5000    |   false   | vrf-device-created | provisioned | cfg-in-sync |
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+--------------+

efa tenant epg show -detail
================================================================================================
=========================================================================
Name          : ten1epg1
Tenant        : ten1
Type          : extension
State         :
Ports         : 10.20.246.15[0/1]
POs           :
Port Property : SwitchPort Mode     : trunk
              : Native Vlan Tagging : false
NW Policy     : Ctag Range          : 11
              : VRF                 : ten1vrf1
              : L3Vni               : 8192
+------+-------------------------+-------+---------+--------------+--------------
+----------------------+---------+----------------+--------------+-------------
+-------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6
|       Local IP        | IPv6 ND |    IPv6 ND     |   IPv6 ND    |  Dev State  |  App
State  |
|      |       Description       |       |         |              |              |
[Device-IP->Local-IP] |   Mtu   | Managed Config | Other Config |
|             |
```

```
+------+-------------------------+-------+---------+--------------+--------------
+-----------------------+---------+----------------+--------------+-------------
+--------------+
|  11  | Tenant L3 Extended VLAN |  11   |         | 10.0.11.1/24 |
|                       |         |     false      |    false     | provisioned | cfg-
in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+-----------------------+---------+----------------+--------------+-------------
+--------------+
================================================================================================
================================================================================================
==
```

2. Complete the following configuration on SLX device:

```
L1# show running-config bridge-
domain 4095
bridge-domain 4095 p2mp
 description Tenant L3 Extended BR
BD
 pw-profile Tenant-profile
 router-interface Ve 8191
!

L1# show running-config interface
Ve 8191
interface Ve 8191
 vrf forwarding ten1vrf1
 ip address 10.40.40.252/31
 ipv6 address
fd40:4040:4040:1::fe/127
 no shutdown
!

L1# do show running-config router
bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.3 remote-as
4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.203/32
  network 172.31.254.226/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  local-as 5000
  redistribute connected
  neighbor 10.40.40.253 remote-as
5000
  neighbor 10.40.40.253 next-hop-
self
  maximum-paths 8
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  neighbor fd40:4040:4040:1::ff
remote-as 5000
  neighbor fd40:4040:4040:1::ff
next-hop-self
  neighbor fd40:4040:4040:1::ff
activate
  maximum-paths 8
 !
!
```

```
L1# show running-config bridge-
domain 4095
bridge-domain 4095 p2mp
 description Tenant L3 Extended BR
BD
 pw-profile Tenant-profile
 router-interface Ve 8191
!

L1# show running-config interface
Ve 8191
interface Ve 8191
 vrf forwarding ten1vrf1
 ip address 10.40.40.252/31
 ipv6 address
fd40:4040:4040:1::fe/127
 no shutdown
!

L1# do show running-config router
bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.3 remote-as
4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.203/32
  network 172.31.254.226/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  local-as 5000
  redistribute connected
  neighbor 10.40.40.253 remote-as
5000
  neighbor 10.40.40.253 next-hop-
self
  maximum-paths 8
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  neighbor fd40:4040:4040:1::ff
remote-as 5000
  neighbor fd40:4040:4040:1::ff
next-hop-self
  neighbor fd40:4040:4040:1::ff
activate
  maximum-paths 8
 !
!
```

Configure Local ASN During VRF Update

You can configure a local ASN when you update a VRF.

About This Task

Follow this procedure to configure a local ASN.

Procedure

1. To configure a local ASN when you update a tenant VRF, run the following command:

```
efa tenant vrf update --name ten1vrf1 --tenant ten1 --operation local-asn-add --local-
asn 6000
WARNING : This operation will result in the reset of the backup routing bgp neighbors
of the VRF. Do you want to proceed (Y/n)?

efa tenant vrf show
+----------+--------+-------------+--------------------+-------------+----------
+-----------+-----------+-------------------+-------------+-------------+
|   Name   | Tenant | Routing Type | Centralized Routers | Redistribute
| Max Path | Local Asn | Enable GR |       State        |  Dev State  |  App State  |
+----------+--------+-------------+--------------------+-------------+----------
+-----------+-----------+-------------------+-------------+-------------+
| ten1vrf1 |  ten1  | distributed  |                     |  connected
|    8     |   6000    |   false   | vrf-device-created | provisioned | cfg-in-sync |
+----------+--------+-------------+--------------------+-------------+----------
+-----------+-----------+-------------------+-------------+-------------+
```

2. Complete the following configuration on SLX device:

```
L1# do show running-config router
bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.3 remote-as
4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.203/32
  network 172.31.254.226/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  local-as 6000
  redistribute connected
  neighbor 10.40.40.253 remote-as
6000
  neighbor 10.40.40.253 next-hop-
self
  maximum-paths 8
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  neighbor fd40:4040:4040:1::ff
remote-as 6000
  neighbor fd40:4040:4040:1::ff
next-hop-self
  neighbor fd40:4040:4040:1::ff
activate
  maximum-paths 8
 !
!
```

```
L2# show running-config router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.2 remote-as
4200000000
 neighbor 10.20.20.2 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.226/32
  network 172.31.254.243/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  local-as 6000
  redistribute connected
  neighbor 10.40.40.252 remote-as
6000
  neighbor 10.40.40.252 next-hop-
self
  maximum-paths 8
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  neighbor fd40:4040:4040:1::fe
remote-as 6000
  neighbor fd40:4040:4040:1::fe
next-hop-self
  neighbor fd40:4040:4040:1::fe
activate
  maximum-paths 8
 !
```

Deconfigure Local ASN during VRF Update

You can deconfigure a local ASN when you update a VRF.

About This Task

Follow this procedure to deconfigure a local ASN.

Procedure

1. To deconfigure a local ASN when you update a tenant VRF, run the following command:

```
efa tenant vrf update --name ten1vrf1 --tenant ten1 --operation local-asn-delete
WARNING : This operation will result in the reset of the backup routing bgp neighbors
of the VRF. Do you want to proceed (Y/n)?


efa tenant vrf show
+--------+------+-----------+------------+------------+----+-----+-------
+------------------+-----------+-----------+
|  Name  |Tenant| Routing   | Centralized|Redistribute|Max
|Local| Enable|      State       | Dev State | App State |
|        |      | Type      | Routers    |
```

```
|Path|ASN  | GR    |                 |           |           |
+--------+------+-----------+------------+-----------+----+-----+-------
+------------------+-----------+-----------+
|ten1vrf1| ten1 |distributed|            | connected  |
8  |     |false  |vrf-device-created|provisioned|cfg-in-sync|
+--------+------+-----------+------------+-----------+----+-----+-------
+------------------+-----------+-----------+
```

2. Complete the following configuration on SLX device:

```
L1# do show running-config router
bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.3 remote-as
4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.203/32
  network 172.31.254.226/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.40.40.253 remote-as
4200000000
  neighbor 10.40.40.253 next-hop-
self
  maximum-paths 8
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  neighbor fd40:4040:4040:1::ff
remote-as 4200000000
  neighbor fd40:4040:4040:1::ff
next-hop-self
  neighbor fd40:4040:4040:1::ff
activate
  maximum-paths 8
 !
!
```

```
L2# show running-config router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.2 remote-as
4200000000
 neighbor 10.20.20.2 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.226/32
  network 172.31.254.243/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.40.40.252 remote-as
4200000000
  neighbor 10.40.40.252 next-hop-
self
  maximum-paths 8
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  neighbor fd40:4040:4040:1::fe
remote-as 4200000000
  neighbor fd40:4040:4040:1::fe
next-hop-self
  neighbor fd40:4040:4040:1::fe
activate
  maximum-paths 8
 !
```

Enable Graceful Restart on Tenant VRF

You can enable graceful restart on each tenant VRF when you create or update a tenant VRF. Based on the endpoints present in the EPG, VRF is instantiated on the

switches when you create an L3 endpoint group or transition L2 endpoint group to L3 endpoint group.

Graceful restart automatically gets configured when you configure a VRF.

About This Task

Follow this procedure to enable graceful restart on tenant VRF.

Note

By default, graceful restart is disabled. Default value of graceful restart is the switch default.

Procedure

1. To enable graceful restart when you create a tenant VRF, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
      --graceful-restart=true/false
```

2. To enable graceful restart on an existing tenant VRF, run the following command:

```
efa tenant vrf update --name <vrf-name> --tenant <tenant-name>
      --operation graceful-restart-update --graceful-restart=true/false
```

   The following example creates a graceful restart on tenant VRF:

```
efa tenant vrf create --name vrf1 --tenant tenant1
efa tenant vrf create -name vrf10 -tenant tenant1 --graceful-restart=true

efa tenant epg create --name ten1epg1 --tenant tenant1
     --port 10.24.80.134[0/11],10.24.80.135[0/11]
     --switchport-mode trunk -ctag-range 11 --vrf vrf1 -anycast-ip 11:10.10.11.1/24
efa tenant epg create --name ten1epg2 --tenant tenant1
     --port 10.24.80.134[0/12],10.24.80.135[0/12]
     --switchport-mode trunk -ctag-range 12 --vrf vrf10 -anycast-ip 12:10.10.12.1/24
efa tenant vrf update --name vrf1 --tenant tenant1
     --operation graceful-restart-update --graceful-restart=true
efa tenant vrf update --name vrf10 --tenant tenant1
     --operation graceful-restart-update -graceful-restart=false
```

Configure MaxPaths on Tenant VRF

XCO allows provisioning of maximum paths for each tenant VRF when you create or update VRF.

About This Task

Follow this procedure to configure maximum paths on tenant VRF.

VRF is updated on the switches when you update a VRF.

Note

- Default value of `max-path` is 8.
- Choosing specific devices for `max-path` provisioning is not allowed.

Procedure

To configure maximum paths when you create a tenant VRF, run the following commands:

```
# efa tenant vrf create --name <vrf-name> --tenant <tenant-name> --max-path <value>

# efa tenant vrf update --name <vrf-name> --tenant <tenant-name> --operation <max-path-
add|max-path-delete> --max-path <value>

# efa tenant vrf create --name vrf1 --tenant tenant1

# efa tenant vrf create -name vrf10 -tenant tenant1

# efa tenant epg create --name ten1epg1 --tenant tenant1 --
port10.24.80.134[0/11],10.24.80.135[0/11] --switchport-mode trunk -ctag-range 11 --vrf
vrf1 --anycast-ip 11:10.10.11.1/24

# efa tenant epg create --name ten1epg2 --tenant tenant1 --
port10.24.80.134[0/12],10.24.80.135[0/12] --switchport-mode trunk -ctag-range 12 --vrf
vrf10 --anycast-ip 12:10.10.12.1/24

# efa tenant vrf update --name vrf10 -tenant tenant1 --operation max-paths-add --max-path
13 Switch Config
```

Configure Resilient Hashing on Tenant VRF

As a load-balancing method, resilient hashing helps to lessen the possibility that a destination path is remapped when a LAG (Link Aggregation Group) link fails.

About This Task

When you create or update a tenant VRF, you can enable ECMP resilient hashing and you can configure the maximum number of resilient hashing paths allowed (8, 16, or 64 paths). Resilient hashing is disabled by default. The default number of allowed paths is the same as the default value for the SLX devices.

Procedure

1. To enable resilient hashing when you create a tenant VRF, use the **`efa tenant vrf create`** command with the **`--rh-ecmp-enable`** and **`--rh-max-path`** parameters, and set the maximum number of allowed paths. For example,

   ```
   # efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
   --rh-ecmp-enable=true/false --rh-max-path <8 | 16 | 64>
   ```

2. To enable resilient hashing on an existing tenant VRF, use the **efa tenant vrf update** command with the **--rh-ecmp-enable** and **--rh-max-path** parameters, and set the maximum number of allowed paths. For example,

```
# efa tenant vrf update --name <vrf-name> --tenant <tenant-name>
--operation <rh-max-path-add | rh-max-path-delete | rh-ecmp-update>
--rh-ecmp-enable=true/false --rh-max-path <8 | 16 | 64>
```

Note

- The **--max-path** and **--rh-max-path** parameters can co-exist.
- You cannot choose the specific devices on which to configure resilient hashing. Configuration applies to all SLX devices in the tenant VRF.
- For more information about the commands, including usage examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Configure Redistribute Attribute on Tenant VRF

XCO allows provisioning of redistribute attribute per tenant VRF when you create or update a VRF.

About This Task

VRF is updated on the switches during the user triggered VRF update operation based on the endpoints present in the endpoint groups using the VRF.

Note

- Default value of `redistribute` is `connected`.
- Choosing specific devices for `redistribute` provisioning is not allowed.

Procedure

1. To configure redistribute attribute when you create a tenant VRF, run the following command:

```
# efa tenant vrf create --name <vrf-name> --tenant <tenant-name> --redistribute <list>
```

2. To configure redistribute attribute on an existing tenant VRF, run the following command:

```
# efa tenant vrf update --name <vrf-name> --tenant <tenant-name> --operation
<redistribute-add|redistribute-delete> --redistribute <static | connected>
```

Example

```
# efa tenant vrf create --name vrf1 --tenant tenant1 --redistribute static

# efa tenant vrf create -name vrf10 -tenant tenant1

# efa tenant epg create --name ten1epg1 --tenant tenant1 --
port10.24.80.134[0/11],10.24.80.135[0/11] --switchport-mode trunk -ctag-range 11 --vrf
vrf1 --anycast-ip 11:10.10.11.1/24

# efa tenant epg create --name ten1epg2 --tenant tenant1 --
port10.24.80.134[0/12],10.24.80.135[0/12] --switchport-mode trunk --ctag-range 12 --
vrf vrf10 --anycast-ip 12:10.10.12.1/24
```

```
# efa tenant vrf update --name vrf10 -tenant tenant1 --operation redistribute-add --
redistribute static
Device1# sh run router bgp
router bgp
local-as 4200000000
address-family ipv4unicast vrf vrf1
redistribute static
!
address-family ipv4unicast vrf vrf10
redistribute connected
redistribute static
!
Device2# sh run router bgp
router bgp
local-as 4200000000
address-family ipv4unicast vrf vrf1
redistribute static
!
address-family ipv4unicast vrf vrf10
redistribute connected
redistribute static
!
address-family ipv6unicast vrf vrf1
redistribute static
!
!address-family ipv6unicast vrf vrf10
redistribute connected
redistribute static
!
```

Configure Advertise Network and Static Network on Tenant VRF

You can configure “network” and “static-network” attributes (advertized by BGP) on a tenant VRF (and device) when you create and update VRF. XCO provisions the “network” and “static-network” attributes on switches when you initiate or update VRF on the switches.

XCO supports only static-network with IPv4.

About This Task

Follow this procedure to configure network and static network advertized by BGP on tenant VRF.

Procedure

1. To configure “network” and “static-network” when you create a tenant VRF, run the following command:

   When you create L3 EPG or transition L2 EPG to L3 EPG, VRF is instantiated on the switches based on the endpoints present in the EPG.

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>

    --ipv4-network <device-ip,network> --ipv4-network-backdoor < device-
ip,network,true|false>
    --ipv4-network-weight <device-ip,network,0-65535> --ipv4-network-route-map <device-
ip,network,route-map>

    --ipv4-static-network <device-ip,static-network> --ipv4-static-network-distance
<device-ip,static-network,1-255>

    --ipv6-network <device-ip,network> --ipv6-network-backdoor <device-ip,network,true|
false>
    --ipv6-network-weight <device-ip,network,0-65535> --ipv6-network-route-map <device-
ip,network,route-map>
```

2. To configure “network” and “static-network” when you update a tenant VRF, run the following command:

   When you trigger VRF update operation, VRF is updated on the switches based on the endpoints present in the EPGs.

```
efa tenant vrf update --name <vrf-name> --tenant <tenant-name>

    --operation network-add|network-delete|static-network-add|static-network-delete

    --ipv4-network <device-ip,network> --ipv4-network-backdoor < device-
ip,network,true|false>
    --ipv4-network-weight <device-ip,network,0-65535> --ipv4-network-route-map <device-
ip,network,route-map>

    --ipv4-static-network <device-ip,static-network> --ipv4-static-network-distance
<device-ip,static-network,1-255>

    --ipv6-network <device-ip,network> --ipv6-network-backdoor <device-ip,network,true|
false>
    --ipv6-network-weight <device-ip,network,0-65535> --ipv6-network-route-map <device-
ip,network,route-map>
```

   The following example configures network and static network (advertized by BGP) on a tenant VRF:

```
efa tenant vrf create --name vrf1 --tenant tenant1
   --ipv4-network 10.24.80.134,10.20.30.40/30
   --ipv4-network 10.24.80.134,10.21.30.40/30 --ipv4-network-backdoor
10.24.80.134,10.21.30.40/30,true

   --ipv4-static-network 10.24.80.134,11.10.30.40/30
   --ipv4-static-network 10.24.80.134,11.20.30.40/30 --ipv4-static-network-distance
10.24.80.134,11.20.30.40/30,169

   --ipv6-network 10.24.80.135,11::22/128
   --ipv6-network 10.24.80.135,11::23/128 --ipv6-network-backdoor
10.24.80.134,11::23/128,true
   --ipv6-network 10.24.80.135,11::24/128 --ipv6-network-weight
10.24.80.134,11::24/128,144
   --ipv6-network 10.24.80.135,11::25/128 --ipv6-network-route-map
10.24.80.134,11::25/128,rmap1
```

```
efa tenant epg create --name ten1epg1 --tenant tenant1
   --port 10.24.80.134[0/11],10.24.80.135[0/11]
   --switchport-mode trunk -ctag-range 11 --vrf vrf1 -anycast-ip 11:10.10.11.1/24

efa tenant vrf update --name vrf1 --tenant tenant1
   --operation network-add
   --ipv4-network 10.24.80.134,10.22.30.40/30 --ipv4-network-weight
10.24.80.134,10.22.30.40/30,144
   --ipv4-network 10.24.80.134,10.23.30.40/30 --ipv4-network-route-map
10.24.80.134,10.23.30.40/30,rmap1
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# sh run router bgp address-
family ipv4 unicast vrf vrf1

router bgp
  address-family ipv4 unicast vrf vrf1
    redistribute connected
    static-network 11.10.30.40/30
    static-network 11.20.30.40/30
distance 169
    network 10.20.30.40/30
    network 10.21.30.40/30 backdoor
    network 10.22.30.40/30 weight 144
    network 10.23.30.40/30 route-map rmap1
  !
  address-family ipv6 unicast vrf vrf1
    redistribute connected
  !
!
```

```
Rack1-Device2# sh run router bgp address-
family ipv4 unicast vrf vrf1

router bgp
  address-family ipv4 unicast vrf vrf1
    redistribute connected
  !
  address-family ipv6 unicast vrf vrf1
    redistribute connected
    network 11::22/128
    network 11::23/128 backdoor
    network 11::24/128 weight 144
    network 11::25/128 route-map rmap1
  !
!
```

Configure Aggregate Address on Tenant VRF

You can configure an aggregate address (advertised by BGP) for ean tenant VRF (and device) when you create or update a VRF. XCO provisions the aggregate address on switches when VRF is instantiated or updated on the switches.

About This Task

Follow this procedure to configure an aggregate address on tenant VRF.

Procedure

1. To configure aggregate address when you create a tenant VRF, run the following command:

   When you trigger L3 EPG create or L2 EPG transition to L3 EPG, VRF is instantiated on the switches based on the endpoints present in the EPG.

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>

     --ipv4-aggregate-address <device-ip,aggregate-address>
     --ipv4-aggregate-summary-only <device-ip,aggregate-address,true|false>
     --ipv4-aggregate-as-set <device-ip,aggregate-address,true|false>
     --ipv4-aggregate-advertise-map <device-ip,aggregate-address,route-map>
     --ipv4-aggregate-suppress-map <device-ip,aggregate-address,route-map>
```

2. To configure aggregate address when you update a tenant VRF, run the following command:

   When you trigger VRF update operation, VRF is updated on the switches based on the endpoints present in the EPGs.

```
efa tenant vrf update --name <vrf-name> --tenant <tenant-name>

     --operation aggregate-address-add| aggregate-address-delete

     --ipv4-aggregate-address <device-ip,aggregate-address>
     --ipv4-aggregate-summary-only <device-ip,aggregate-address,true|false>
     --ipv4-aggregate-as-set <device-ip,aggregate-address,true|false>
     --ipv4-aggregate-advertise-map <device-ip,aggregate-address,route-map>
     --ipv4-aggregate-suppress-map <device-ip,aggregate-address,route-map>
```

   The following example configures aggregate address during VRF create operation:

```
efa tenant vrf create --name vrf1 --tenant tenant1

     --ipv4-aggregate-address 10.24.80.134,10.20.21.40/30
     --ipv4-aggregate-summary-only 10.24.80.134,10.20.21.40/30,true
     --ipv4-aggregate-as-set 10.24.80.134,10.20.21.40/30,true
     --ipv4-aggregate-advertise-map 10.24.80.134,10.20.21.40/30,some
     --ipv4-aggregate-suppress-map 10.24.80.134,10.20.21.40/30,some

     --ipv6-aggregate-address 10.24.80.135,10::20/126
     --ipv6-aggregate-summary-only 10.24.80.135,10::20/126,true
     --ipv6-aggregate-as-set 10.24.80.135,10::20/126,true
     --ipv6-aggregate-advertise-map 10.24.80.135,10::20/126,some
     --ipv6-aggregate-suppress-map 10.24.80.135,10::20/126,some

efa tenant epg create --name ten1epg1 --tenant tenant1
     --port 10.24.80.134[0/11],10.24.80.135[0/11]
     --switchport-mode trunk –ctag-range 11 --vrf vrf1 –anycast-ip 11:10.10.11.1/24

efa tenant vrf update –-name vrf1 –-tenant tenant1
     --operation aggregate-address-add

     --ipv4-aggregate-address 10.24.80.134,10.21.21.40/30
     --ipv4-aggregate-summary-only 10.24.80.134,10.21.21.40/30,true
     --ipv4-aggregate-as-set 10.24.80.134,10.21.21.40/30,true
     --ipv4-aggregate-advertise-map 10.24.80.134,10.21.21.40/30,some
     --ipv4-aggregate-suppress-map 10.24.80.134,10.21.21.40/30,some

     --ipv6-aggregate-address 10.24.80.135,11::20/126
     --ipv6-aggregate-summary-only 10.24.80.135,11::20/126,true
     --ipv6-aggregate-as-set 10.24.80.135,11::20/126,true
     --ipv6-aggregate-advertise-map 10.24.80.135,11::20/126,some
     --ipv6-aggregate-suppress-map 10.24.80.135,11::20/126,some
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# sh run router bgp address-
family ipv4 unicast vrf vrf1

router bgp
  address-family ipv4 unicast vrf vrf1
    redistribute connected
    aggregate-address 10.20.21.40/30
advertise-map some
    aggregate-address 10.20.21.40/30 as-
set
    aggregate-address 10.20.21.40/30
summary-only
```

```
Rack1-Device2# sh run router bgp address-
family ipv4 unicast vrf vrf1

router bgp
  address-family ipv4 unicast vrf vrf1
    redistribute connected
  !
  address-family ipv6 unicast vrf vrf1
    redistribute connected
    aggregate-address 10::20/126
advertise-map some
    aggregate-address 10::20/126 as-set
```

```
    aggregate-address 10.20.21.40/30
suppress-map some
    aggregate-address 10.21.21.40/30
advertise-map some
    aggregate-address 10.21.21.40/30 as-
set
    aggregate-address 10.21.21.40/30
summary-only
    aggregate-address 10.21.21.40/30
suppress-map some
  !
  address-family ipv6 unicast vrf vrf1
    redistribute connected
  !
!
```

```
    aggregate-address 10::20/126 summary-
only
    aggregate-address 10::20/126 suppress-
map some
    aggregate-address 11::20/126
advertise-map some
    aggregate-address 11::20/126 as-set
    aggregate-address 11::20/126 summary-
only
    aggregate-address 11::20/126 suppress-
map some
  !
!
```

Configure EVPN IRB VE Cluster Gateway on a Tenant VRF

You can enable EVPN IRB VE cluster gateway.

About This Task

Follow this procedure to configure an EVPN IRB VE Cluster-gateway on tenant VRF.

Note

- A layer3-extension is enabled by default for a Distributed VRF instantiated by XCO and you cannot disable it.
- A layer3-extension is disabled by default for a Centralized VRF instantiated by XCO and you cannot enable it.

When a layer3-extension is enabled, XCO pushes the following configurations:

1. EVPN IRB BD
2. EVPN IRB VE
3. Addition of the EVPN IRB BD to the EVPN instance
4. Addition of the EVPN IRB BD to the overlay-gateway instance
5. EVPN IRB VE configuration under the VRF

Procedure

1. To configure EVPN IRB VE cluster gateway on a distributed tenant VRF, Run the following command:

```
efa tenant vrf create --tenant <tenant-name> --name <vrf-name>
           --layer3-extension-enable {true | false}

efa tenant vrf create --tenant "t1" --name "v1" --routing-type "distributed" --rt-type
import --rt 101:101 --rt-type export --rt 101:101

efa tenant vrf show --name v1 --tenant t1 -detail
================================================================================
Name                          : v1
Tenant:                       : t1
Routing Type                  : distributed
Centralized Routers           :
Enable Layer3 Extension       : true
Redistribute                  : connected
Max Path                      : 8
```

```
Local Asn                   :
L3VNI                       :
EVPN IRB BD                 :
EVPN IRB VE                 :
BR VNI                      :
BR BD                       :
BR VE                       :
RH Max Path                 :
Enable RH ECMP              : false
Enable Graceful Restart     : false
Route Target                :
Static Route                :
Static Route BFD            :
Network Route Address       :
Static Network              :
Aggregate Address           :
VRF Type                    :
State                       : vrf-created
Dev State                   : not-provisioned
App State                   : cfg-ready
================================================================================


efa tenant epg create --name epg1 --tenant t1 --switchport-mode trunk –port
10.20.246.15[0/1] --vrf v1 --switchport-native-vlan 10 --ctag-range 10 --anycast-ip
10:10.10.12.1/24


efa tenant vrf show --name v1 --tenant t1 –detail
================================================================================
Name                        : v1
Tenant                      : t1
Routing Type                : distributed
Centralized Routers         :
Enable Layer3 Extension     : true
Redistribute                : connected
Max Path                    : 8
Local Asn                   :
L3VNI                       : 10111
EVPN IRB BD                 : 4096
EVPN IRB VE                 : 8192
BR VNI                      : 10110
BR BD                       : 4095
BR VE                       : 8191
RH Max Path                 :
Enable RH ECMP         : false
Enable Graceful Restart     : false
Route Target                : import 101:101
                            : export 101:101
Static Route           :
Static Route BFD            :
Network Route Address       :
Static Network         :
Aggregate Address           :
VRF Type                    :
State                       : vrf-device-created
Dev State                   : provisioned
```

```
App State                  : cfg-in-sync
================================================================================
```

a. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
vrf v1
vrf v1
 rd 172.31.254.19:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
 address-family ipv6 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
!
```

```
Rack1-Device2# show running-config
vrf v1
vrf v1
 rd 172.31.254.20:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
 address-family ipv6 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
!
```

2. To configure EVPN IRB VE cluster-gateway on a centralized tenant VRF without layer3-extension, run the following command:

```
efa tenant vrf create --tenant <tenant-name> --name <vrf-name>
                      --layer3-extension-enable {true | false}
efa tenant vrf create --tenant "t1" --name "v2" --routing-type "centralized"

efa tenant vrf show --name v2 --tenant t1 -detail
================================================================================
Name                       : v2
Tenant                     : t1
Routing Type               : centralized
Centralized Routers        : 10.20.246.15
                           : 10.20.246.16
Enable Layer3 Extension    : false
Redistribute               : connected
Max Path                   : 8
Local Asn                  :
L3VNI                      :
EVPN IRB BD                :
EVPN IRB VE                :
BR VNI                     :
BR BD                      :
BR VE                      :
RH Max Path                :
Enable RH ECMP             : false
Enable Graceful Restart    : false
Route Target               :
Static Route               :
Static Route BFD           :
Network Route Address      :
Static Network             :
Aggregate Address          :
VRF Type                   :
State                      : vrf-created
Dev State                  : not-provisioned|
App State                  : cfg-ready
================================================================================

efa tenant epg create --name epg1 --tenant t1 --switchport-mode trunk -port
10.20.246.15[0/1] --vrf v2 --switchport-native-vlan 10 --ctag-range 10 --anycast-ip
10:10.10.12.1/24
```

```
efa tenant vrf show --name v2 --tenant t1 -detail
================================================================================
Name                        : v2
Tenant                      : t1
Routing Type                : centralized
Centralized Routers         : 10.20.246.15
                            : 10.20.246.16
Enable Layer3 Extension     : false
Redistribute                : connected
Max Path                    : 8
Local Asn                   :
L3VNI                       :
EVPN IRB BD                 :
EVPN IRB VE                 :
BR VNI                      : 10110
BR BD                       : 4096
BR VE                       : 8192
RH Max Path                 :
Enable RH ECMP              : false
Enable Graceful Restart     : false
Route Target                : import 101:101
                            : export 101:101
Static Route          :
Static Route BFD            :
Network Route Address       :
Static Network        :
Aggregate Address           :
VRF Type                    :
State                       : vrf-device-created
Dev State                   : provisioned
App State                   : cfg-in-sync
================================================================================
=
```

3. To configure EVPN IRB VE cluster gateway on a centralized tenant VRF with layer3-extension, Run the following command:

```
efa tenant vrf create --tenant <tenant-name> --name <vrf-name>
             --layer3-extension-enable {true | false}
efa tenant vrf create --tenant "t1" --name "v3" --routing-type "centralized" --layer3-
extension-enable true

efa tenant vrf show --name v3 --tenant t1 -detail
================================================================================
Name                        : v3
Tenant                      : t1
Routing Type                : centralized
Centralized Routers         : 10.20.246.15
                            : 10.20.246.16
Enable Layer3 Extension     : true
Redistribute                : connected
Max Path                    : 8
Local Asn                   :
L3VNI                       :
EVPN IRB BD                 :
EVPN IRB VE                 :
BR VNI                      :
BR BD                       :
BR VE                       :
RH Max Path                 :
Enable RH ECMP              : false
Enable Graceful Restart     : false
Route Target                :
Static Route                :
Static Route BFD            :
```

```
Network Route Address     :
Static Network            :
Aggregate Address         :
VRF Type                  :
State                     : vrf-created
Dev State                 : not-provisioned
App State                 : cfg-ready
================================================================================

efa tenant epg create --name epg1 --tenant t1 --switchport-mode trunk -port
10.20.246.15[0/1] --vrf v3 --switchport-native-vlan 10 --ctag-range 10 --anycast-ip
10:10.10.12.1/24


efa tenant vrf show --name v3 --tenant t1 -detail
================================================================================
Name                      : v3
Tenant                    : t1
Routing Type              : centralized
Centralized Routers       : 10.20.246.15
                          : 10.20.246.16
Enable Layer3 Extension   : true
Redistribute              : connected
Max Path                  : 8
Local Asn                 :
L3VNI                     : 10111
EVPN IRB BD               : 4096
EVPN IRB VE               : 8192
BR VNI                    : 10110
BR BD                     : 4095
BR VE                     : 8191
RH Max Path               :
Enable RH ECMP            : false
Enable Graceful Restart   : false
Route Target              : import 101:101
                          : export 101:101
Static Route              :
Static Route BFD          :
Network Route Address     :
Static Network            :
Aggregate Address         :
VRF Type                  :
State                     : vrf-device-created
Dev State                 : provisioned
App State                 : cfg-in-sync
================================================================================
```

a. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
vrf v3
vrf v2
 rd 172.31.254.19:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
 address-family ipv6 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
!
```

```
Rack1-Device2# show running-config
vrf v3
vrf v2
 rd 172.31.254.20:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
 address-family ipv6 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
 !
!
```

Route Distinguisher (RD) Allocation Independent of Route Target (RT)

You can allocate route distinguisher which is independent of route target.

Provisioning in EFA 2.5.5 or above

The following example shows the allocation of route distinguisher:

Note

- XCO auto allocates VRF RT in the format xx:**yy** and VRF RD in the format <router-id>:<**unique-number-per-vrf-independent-of-rt**>.
- RD value has no relation to yy value of RT. XCO auto allocates an unique number of RD for each VRF which is appended to the router ID.

```
efa tenant vrf create --tenant "ten1" --name "ten1vrf1" --routing-type "distributed" --
rt-type export --rt 65010:1 --rt-type import --rt 65010:2 --max-path 8 --redistribute
connected

efa tenant vrf create --tenant "ten1" --name "ten1vrf2" --routing-type "distributed" --
rt-type export --rt 65010:2 --rt-type import --rt 65010:1 --max-path 8 --redistribute
connected

efa tenant epg create --tenant "ten1" --name "ten1epg1" --type extension --switchport-
mode trunk --single-homed-bfd-session-type auto --po ten1po1 --vrf ten1vrf1 --ctag-range
25 --l3-vni 32821 --anycast-ip 25:10.0.21.1/24 --ctag-description "25:Tenant L3 Extended
VLAN" --l2-vni 25:32770 --suppress-arp 25:true --suppress-nd 25:false

efa tenant epg create --tenant "ten1" --name "ten1epg2" --type extension --switchport-
mode trunk --single-homed-bfd-session-type auto --po ten1po1 --vrf ten1vrf2 --ctag-range
26 --l3-vni 32823 --anycast-ip 26:11.0.21.1/24 --ctag-description "26:Tenant L3 Extended
VLAN" --l2-vni 26:32771 --suppress-arp 26:true --suppress-nd 26:false
```

Switch Config in EFA 2.5.5 or above

Verify the following switch configuration on SLX devices:

```
Rack1-Device1# show running-config
vrf
vrf ten1vrf1
rd 172.31.254.40:1
evpn irb ve 8192
address-family ipv4 unicast
  route-target export 65010:1 evpn
  route-target import 65010:2 evpn
!
address-family ipv6 unicast
  route-target export 65010:1 evpn
  route-target import 65010:2 evpn
!
!
vrf ten1vrf2
rd 172.31.254.40:2
evpn irb ve 8190
address-family ipv4 unicast
  route-target export 65010:2 evpn
  route-target import 65010:1 evpn
!
address-family ipv6 unicast
  route-target export 65010:2 evpn
  route-target import 65010:1 evpn
!
```

```
Rack1-Device2# show running-config
vrf
vrf ten1vrf1
rd 172.31.254.209:1
evpn irb ve 8192
address-family ipv4 unicast
  route-target export 65010:1 evpn
  route-target import 65010:2 evpn
!
address-family ipv6 unicast
  route-target export 65010:1 evpn
  route-target import 65010:2 evpn
!
!
vrf ten1vrf2
rd 172.31.254.209:2
evpn irb ve 8190
address-family ipv4 unicast
  route-target export 65010:2 evpn
  route-target import 65010:1 evpn
!
address-family ipv6 unicast
  route-target export 65010:2 evpn
  route-target import 65010:1 evpn
!
```

IPv6 Anycast Gateway Support

XCO supports the provisioning of IPv6 anycast gateways.

The following commands support provisioning of anycast gateways in XCO.

```
efa tenant vrf create --name vrf1 --tenant tenant1 --rt-type both --rt 1:1
efa tenant epg create --name ten1epg1 --tenant tenant1 --port 10.24.80.134[0/15]
--switchport-mode trunk --ctag-range 1001 --anycast-ip 1001:10.0.1.1/24
--anycast-ipv6 1001:2001:10:0:1::1/64 --vrf vrf1
```

Configure Static VRF Route

The static route configuration at the tenant VRF level enables you to provide static routes for each tenant VRF.

About This Task

Follow this procedure to configure static VRF route on a tenant VRF.

Procedure

1. To configure static VRF route, run the following commands:

```
# efa tenant vrf create [ --name vrf name | --tenant tenant name | --rttype
<both | import | export > | --rt | --ipv4-static-route-next-hop < device
ip,ipv4 static route network,nexthop ip,nexthop distance metric> | --ipv6-static-
route-next-hop < device ip,ipv6 static route network,nexthop ip,nexthop distance
metric> | --local-asn | --ipv4-static-route-bfd < device-ip,dest-ipv4-addr,source-
ipv4-addr[interval,minrx,multiplier] > | --ipv6-static-route-bfd < device-ip,dest-ipv4-
addr,source-ipv4-addr[interval,min-rx,multiplier] > | --max-path uint <1-64> | --
redistribute < static | connected >| --rh-max-path uint | --rh-ecmp-enable | --help ]
```

2. To update static VRF route, run the following commands:

```
# efa tenant vrf update [ --name <vrf-name> --tenant <tenant-name> --operation
< staticroute-add | static-route-delete> --ipv6static-route-next-hop <destination,
next-hop> --ipv4static-route-next-hop <destination, next-hop> efa tenant vrf update
[ --name vrf name | --tenant tenant name | --operation < local-asn-add |
local-asn-delete | static-route-bfd-add | static-route-bfd-delete | static-route-
add | static-route-delete | max-path-add | max-path-delete | redistribute-add |
redistributedelete | rh-max-path-add | rh-max-path-delete | rh-ecmp-update > | --local-
asn | --ipv4-static-route-bfd < device IP,dest-ipv4-addr,source-ipv4-addr[interval,min-
rx,multiplielr] > | --ipv6-staticroute-bfd < device ip,dest-ipv6-addr,source-
ipv6-addr[interval,minrx,multiplielr] > | --ipv4-static-route-next-hop < device-
ip,destipv4- addr,source-ipv4-addr[interval,min-rx,multiplier],distance,metric > |
--ipv6-static-route-next-hop < device-ip,dest-ipv6-addr,source-ipv6-addr[interval,min-
rx,multiplier],distance,metric > | --max-path uint <1-64> | --redistribute < static |
connected > | --rh-max-path uint | --rh-ecmpenable | --help ]
```

The following example configures static VRF route:

```
# efa tenant vrf create --name red --tenant tenant11 --ipv6-static-route-next-hop
10.24.80.134,2000::/64,1001::2,,2 --ipv6-static-route-next-hop
10.24.80.134,2000::/64,1002::2 --ipv6-static-route-next-hop
10.24.80.134,2000::/64,1003::2,,4 --ipv6-static-route-next-hop
10.24.80.134,2000::/64,1004::2 --ipv6-static-route-next-hop
10.24.80.135,2001::/64,1001::2,4 --ipv6-static-route-next-hop
10.24.80.135,2001::/64,1002::2 --ipv6-static-route-next-hop
10.24.80.135,2001::/64,1003::2 --ipv6-static-route-next-hop
10.24.80.135,2001::/64,1004::2 --ipv4-static-route-next-hop
10.24.80.134,22.0.0.0/24,13.0.0.1,2,9 --ipv4-static-route-next-hop
10.24.80.134,22.0.0.0/24,13.0.0.2,,7 --ipv4-static-route-next-hop
10.24.80.134,22.0.0.0/24,13.0.0.3 --ipv4-static-route-next-hop
10.24.80.134,22.0.0.0/24,13.0.0.4 --ipv4-static-route-next-hop
10.24.80.135,23.0.0.0/24,13.0.0.1 --ipv4-static-route-next-hop
10.24.80.135,23.0.0.0/24,13.0.0.2 --ipv4-static-route-next-hop
10.24.80.135,23.0.0.0/24,13.0.0.3 --ipv4-static-route-next-hop
10.24.80.135,23.0.0.0/24,13.0.0.4
# efa tenant vrf show --name red --tenant tenant11
================================================================================
Name                : red
Tenant Name         : tenant11
L3 VNI              :
Route Target        :
Static Route        : Switch-IP->{Network,Nexthop-IP[Route-Distance,Route-Metric]}, ...
                    : 10.24.80.134->{22.0.0.0/24,13.0.0.1[2,9]}
{22.0.0.0/24,13.0.0.2[ ,7]}  {22.0.0.0/24,13.0.0.3[ , ]}  {22.0.0.0/24,13.0.0.4[ , ]}
{2000::/64,1001::2[ ,2]}
                           {2000::/64,1002::2[ , ]}  {2000::/64,1003::2[ ,4]}
{2000::/64,1004::2[ , ]}
                    : 10.24.80.135->{23.0.0.0/24,13.0.0.1[ , ]}
{23.0.0.0/24,13.0.0.2[ , ]}  {23.0.0.0/24,13.0.0.3[ , ]}  {23.0.0.0/24,13.0.0.4[ , ]}
{2001::/64,1001::2[4, ]}
                           {2001::/64,1002::2[ , ]}  {2001::/64,1003::2[ , ]}
{2001::/64,1004::2[ , ]}
Local Asn :
================================================================================
# efa tenant epg create --name ten1epg1 --tenant tenant1 --port
10.24.80.134[0/11],10.24.80.135[0/11] --switchport-mode trunk --ctag-range 11 --vrf red
--anycast-ip 11:10.10.11.1/24
```

The `metric` attribute in EFA 3.0.0 and above has a value range from 1 through 16 supported by SLX.

Configure BFD on Static VRF Route

At the tenant VRF level, Bidirectional Forwarding Detection (BFD) configuration enables you to use static route BFD timers.

About This Task

Follow this procedure to configure BFD on static VRF route.

Note

For more information on BFD timers, see BFD Timers for Router BGP BFD and Static Route BFD Sessions on page 222.

Procedure

1. To configure BFD when you create a tenant VRF, run the following command:

```
# efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
--ipv6static-route-bfd <destination-ip, source-ip, bfd-min-tx, bfd-min-rx, bfd-
multiplier>
--ipv4static-route-bfd < destination-ip, source-ip, bfd-min-tx, bfd-min-rx, bfd-
multiplier>
```

2. To configure BFD on an existing tenant VRF, run the following command:

```
# efa tenant vrf update -name <vrf-name> --tenant <tenant-name> --operation <static-
route-bfd-add|static-route-bfd-delete> --ipv6-static-route-bfd <switch-ip, destination-
ip, source-ip, bfd-min-tx, bfd-min-rx, bfd-multiplier> --ipv4static-route-bfd  <switch-
ip, destination-ip, source-ip, bfd-min-tx, bfd-min-rx, bfd-multiplier>
```

Example

The following example configures BFD on static VRF route:

```
# efa tenant vrf create --name red --tenant
tenant11 --ipv6-static-route-bfd 10.24.80.134,1001::2,1001::1,100,200,5 --ipv6-
static-route-bfd 10.24.80.135,1011::2,1011::1,100,200,5 --ipv6-static-route-bfd
10.24.80.134,1002::2, 1002::1 --ipv6-static-route-bfd 10.24.80.135,1012::2,
1012::1 --ipv4-static-route-bfd 10.24.80.134,13.0.0.1,13.0.0.9,200,300,6 --ipv4-
static-route-bfd 10.24.80.135,13.0.1.1,13.0.1.9,200,300,6 --ipv4-static-route-bfd
10.24.80.134,13.0.0.2,13.0.0.10 --ipv4-static-route-bfd 10.24.80.135,13.0.1.2,13.0.1.10

# efa tenant epg create --name ten1epg1 --tenant tenant1 --port
10.24.80.134[0/11],10.24.80.135[0/11] --switchport-mode trunk -ctag-range 11 --vrf red -
anycast-ip 11:10.10.11.1/24
```

Configure Backup Routing on Tenant VRF

You can enable backup routing when all the links from a leaf device to the spine layer are down, and tenant traffic is to be routed via the MCT neighbor.

About This Task

Follow this procedure to configure backup routing on tenant VRF.

A pair of IPv4 and IPv6 address is allocated to each MCT pair across all the tenant VRFs, and the BGP session is established between the same IP Pair.

Procedure

1. Allocate a pair of IPv4 and IPv6 address to each MCT pair across all the tenant VRFs.
   a. Allocate a bridge domain per VRF for backup routing.
   b. Allocate a corresponding router-interface VE per BD or VRF.
   c. Assign the IPv4 and IPv6 address allocated to each device on each of the VE interface.

      Note
      Same IPv4 and IPv6 address is allocated on each of the VE interface which belongs to different VRF.

   d. Establish IBGP IPv4 neighborship with the MCT peer on a set of IP address per VRF.
   e. Establish IBGP IPv6 neighborship with the MCT peer on a set of IPv6 address per VRF.
   f. Configure "next-hop-self" on both the IPv4 and IPv6 neighbor.
   g. Configure "active" on the IPv6 neighbor.

   Example:

```
efa fabric setting update --name fabric1
--backup-routing-ipv4-range 21.1.1.0/24 --backup-routing-ipv6-range 2001:21:1:1::0/120
```

   Example when backup routing is enabled:

```
efa fabric setting update --name nc --backup-routing-enable yes
```

2. Configure devices on tenant VRFs.

   The following table provides an example of device configuration on a tenant VRF:

Table 16: Tenant1 VRF “vrf1”

leaf-9250-173# show running-config
bridge-domain 3001
bridge-domain 3001 p2mp
 pw-profile default
 router-interface Ve 7001
 bpdu-drop-enable
!
 	
leaf-9250-173# sh run in ve 7001
interface Ve 7001
 vrf forwarding vrf1
 ip address 21.1.1.10/31
 ipv6 address 2001:21:1:1::10/127
 no shutdown
!
	leaf-9250-173# show running-config
	router bgp address-family ipv4
	unicast vrf vrf1
	router bgp
	 address-family ipv4 unicast vrf
	vrf1
	  local-as 4210000001
	  redistribute connected
	  neighbor 21.1.1.11 remote-as
	4210000001
	  neighbor 21.1.1.11 next-hop-self
	  maximum-paths 2
	 !
	!
	
	leaf-9250-173# show running-config
	router bgp address-family ipv6
	unicast vrf vrf1
	router bgp
	 address-family ipv6 unicast vrf
	vrf1
	  redistribute connected
	  neighbor 2001:21:1:1::11 next-
	hop-self
	  neighbor 2001:21:1:1::11 remote-
	as 4210000001
	  neighbor 2001:21:1:1::11 activate
	  maximum-paths 2
	 !
	!

Table 17: Tenant2 VRF “vrf2”

leaf-9250-173# show running-config
bridge-domain 3002
bridge-domain 3002 p2mp
 pw-profile default
 router-interface Ve 7002
 bpdu-drop-enable
!
 	
leaf-9250-173# sh run in ve 7002
interface Ve 7002
 vrf forwarding vrf2
 ip address 21.1.1.10/31
 ipv6 address 2001:21:1:1::10/127
 no shutdown
!
	leaf-9250-173# show running-config
	router bgp address-family ipv4
	unicast vrf vrf2
	router bgp
	 address-family ipv4 unicast vrf
	vrf2
	  local-as 4210000001
	  redistribute connected
	  neighbor 21.1.1.11 remote-as
	4210000001
	  neighbor 21.1.1.11 next-hop-self
	  maximum-paths 2
	 !
	!
	
	leaf-9250-173# show running-config
	router bgp address-family ipv6
	unicast vrf vrf2
	router bgp
	 address-family ipv6 unicast vrf
	vrf2
	  redistribute connected

Table 17: Tenant2 VRF “vrf2” (continued)

	neighbor 2001:21:1:1::11 next-hop-self neighbor 2001:21:1:1::11 remote-as 4210000001 neighbor 2001:21:1:1::11 activate maximum-paths 2 ! !

Distributed and Centralized Routing

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.

You can provide centralized mode or distributed mode as input when you configure a router.

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

The default option is centralized

The L3 service plugin passes the routing information to XCO. Based on the configuration mode, the XCO tenant service configures VRF or routing on border leafs or on leaf nodes.

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. The creation of a distributed router is disabled.

Prepare Clos Fabric for Centralized Routing

You can use the following command to create a Clos fabric containing border-leaf devices, which can be used for centralized routing.

Tip

If any device in a fabric is in "admin-down" state, the following commands in the same fabric will not add or delete devices: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

```
efa fabric create --name <fabric-name> --type clos

efa fabric device add-bulk --name <fabric-name>
      --border-leaf <list-of-border-leaf-ip> --leaf <list-of-leaf-ip>
      --spine <list-of-spine-ip>

efa fabric configure --name <fabric-name>
```

Prepare Small Data Center Fabric for Centralized Routing

You can use the following command to create a small data center (non-Clos) fabric containing border-leaf devices, which can be used for centralized routing.

Tip
If any device in a fabric is in "admin-down" state, the following commands in the same fabric will not add or delete devices: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

```
efa fabric create --name <fabric-name> --type non-clos

efa fabric device add-bulk --name <fabric-name>
           --rack <leaf-rack-name> --ip <leaf-ip-pair>
           --border-leaf-rack <bl-rack-name> --border-leaf-ip <bl-ip-pair>

efa fabric configure --name <fabric-name>
```

Enable Centralized Routing on Tenant VRF

You can enable centralized routing at the tenant VRF level to override the default distributed routing behavior. A given tenant can have multiple VRFs with some VRFs operating in distributed routing mode and some VRFs operating in centralized routing mode.

About This Task

Follow this procedure to enable centralized routing.

Procedure

To enable centralized routing, run the following command when you create a tanant VRF:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
    --routing-type {centralized | distributed}
    --centralized-router <list-of-border-leaf-routers>
```

Example

The following example enables centralized routing on a tenant VRF:

```
efa tenant vrf create --name VRF1 --tenant tenant1
    --routing-type centralized --centralized-router BL1-IP,BL2-IP
```

Configure Physical Router for Centralized Routing on Tenant VRF

You can configure a physical router for centralized routing.

About This Task

Follow this procedure to configure a physical router for centralized routing.

Note

- Provide a list of border leaf IPs on which the centralized router (VRF) needs to be added. If a given fabric has only one BL pair and you have not provided any BL pair as centralized router, then the only available BL pair is used as centralized router by default.
- Provide only one BL pair on which the centralized router (VRF) needs to be added.
- VRF instantiation happens on the border leaf devices during the EPG (endpoint group) create or update operations.
- You cannot provide leaf, spine, or super-spine IPs as the target device for centralized routing.
- VRF (with centralized routing enabled) and its dependent L3 configuration (anycast-ip, local-ip, VRF static route, VRF static route bfd, router bgp static or dynamic peer, and router bgp peer-group) are instantiated only on the border leafs on which the parent VRF exists.

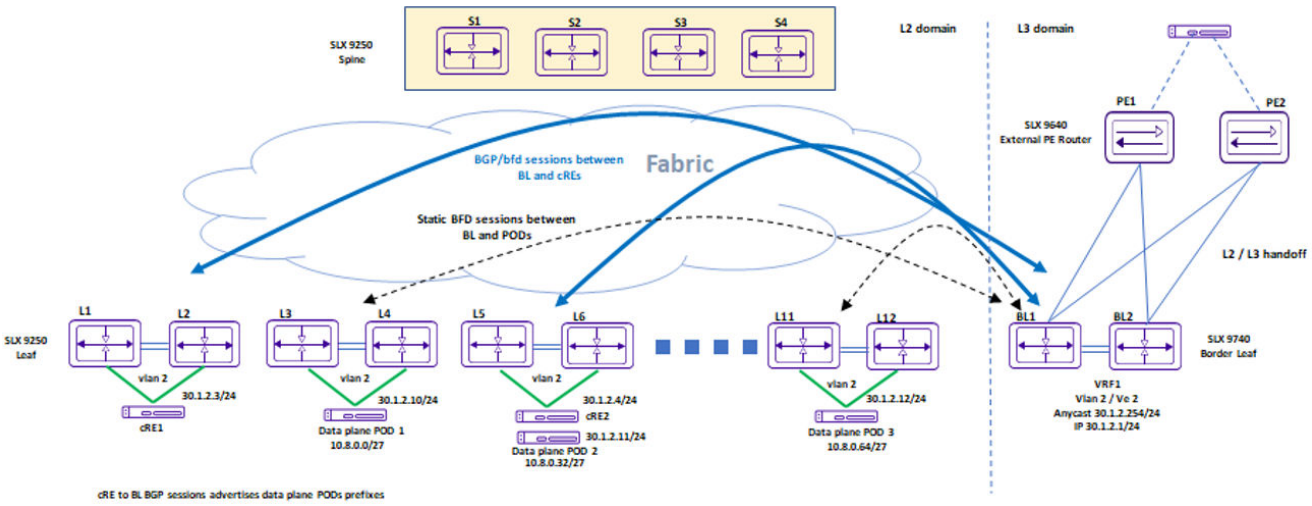

When a centralized routing is enabled for a given tenant VRF:

1. Define a target border leaf device on which the VRF needs to be instantiated.

   The VRF instantiation happens only on those border leaf devices and not on any other leaf or border leaf devices.
2. You do **not** need to provide a target border leaf device on which the anycast IP needs to be configured.

   The anycast IP is configured *automatically* on the border leaf devices on which the VRF is instantiated.

3. Provide a border leaf IP (on which the VRF is instantiated) for the local IP configuration.
4. Provide a border leaf IP (on which the VRF is instantiated) for the VRF SR (Static Route) and VRF SR-BFD (Static Route – BFD) configuration.
5. Provide a border leaf IP (on which the VRF is instantiated) for the BGP static and dynamic peer configuration.
6. Provide a border leaf IP (on which the VRF is instantiated) for the BGP peer-group configuration.

Procedure

1. To configure physical routers for centralized routing on Tenant VRF, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
    --routing-type {centralized | distributed}
    --centralized-router <list-of-border-leaf-routers>
```

   Example

```
efa tenant vrf create --name VRF1 --tenant tenant1
    --routing-type centralized --centralized-router BL1-IP,BL2-IP
```

2. Carve out of VRFs on the border-leaf pairs.

   Instantiate VRFs on the border-leaf devices based on the L3 scale requirements.

   Suppose the fabric has 100 VRFs with 4K anycast IP, then you can instantiate all 100 VRFs on a single border-leaf pair. If the L3 scale requirements are higher than the scale supported by a single border-leaf pair, then add additional border-leaf pair.

Configure Anycast IP on Tenant Endpoint Group

Anycast IP automatically gets configured on the border leaf devices (BL1 and BL2) on which the VRF is instantiated.

```
efa tenant epg create --name ten1epg1 --tenant tenant1
    --port L1-IP[0/11],L2-IP[0/11]
    --switchport-mode trunk –ctag-range 11 --vrf VRF1 --anycast-ip 11:10.10.11.1/24
```

Configure Local IP on Tenant Endpoint Group

You can configure a local IP on tenant endpoint group.

About This Task

Follow this procedure to configure a local IP.

Procedure

To configure a local IP on tenant endpoint group (EPG), run the following commands:

```
efa tenant epg create --name ten1epg1 --tenant tenant1 --vrf VRF1 --switchport-mode trunk
--ctag-range 11
    --anycast-ip 11:10.10.11.1/24 --port L1-IP[0/1],L2-IP[0/1]
    --local-ip 11,BL1-IP:11.22.33.41/24 --local-ip 11,BL2-IP:11.22.34.41/24
```

Configure Static Route on Tenant VRF

Provide the border-leaf IP (on which the VRF is instantiated) for the VRF SR (Static Route) and VRF SR-BFD (Static Route – BFD) configuration.

About This Task

Follow this procedure to configure a static route on tenant VRF.

Procedure

1. To create a static route on tenant VRF, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
    --ipv6-static-route-next-hop <border-leaf-ip, destination, next-hop, distance,
metric>
    --ipv4-static-route-next-hop <border-leaf-ip, destination, next-hop, distance,
metric>
```

2. To update a static route on tenant VRF, run the following command:

```
efa tenant vrf update –name <vrf-name> --tenant <tenant-name>
    --operation <static-route-add|static-route-delete>
    --ipv6-static-route-next-hop <border-leaf-ip, destination, next-hop, distance>
    --ipv4-static-route-next-hop <border-leaf-ip, destination, next-hop, distance>
```

Example

The following example creates static routes on tenant VRF:

```
efa tenant vrf create --name VRF1 --tenant tenant1
    --ipv6-static-route-next-hop BL1-IP,2000::/64,1001::2
    --ipv6-static-route-next-hop BL1-IP,2000::/64,1002::2
    --ipv6-static-route-next-hop BL2-IP,2001::/64,1001::2,4
    --ipv6-static-route-next-hop BL2-IP,2001::/64,1002::2
    --ipv4-static-route-next-hop BL1-IP,22.0.0.0/24,13.0.0.1,2
    --ipv4-static-route-next-hop BL1-IP,22.0.0.0/24,13.0.0.2
    --ipv4-static-route-next-hop BL2-IP,23.0.0.0/24,13.0.0.1
    --ipv4-static-route-next-hop BL2-IP,23.0.0.0/24,13.0.0.2
```

Configure Static Route BFD on Tenant VRF

Provide a border-leaf IP (on which the VRF is instantiated) for the VRF SR-BFD (Static Route – BFD) configuration.

About This Task

Follow this procedure to configure a static route BFD on tenant VRF.

Procedure

1. To configure static route BFD on Tenant VRF when you create a VRF, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
    --ipv6-static-route-bfd <border-leaf-ip, destination-ip, source-ip, bfd-min-tx,
bfd-min-rx, bfd-multiplier>
    --ipv4-static-route-bfd <border-leaf-ip, destination-ip, source-ip, bfd-min-tx,
bfd-min-rx, bfd-multiplier>
```

2. To configure static route BFD on Tenant VRF when you create a VRF, run the following command:

```
efa tenant vrf update –name <vrf-name> --tenant <tenant-name>
    --operation <static-route-bfd-add|static-route-bfd-delete>
    --ipv6-static-route-bfd <border-leaf-ip, destination-ip, source-ip, bfd-min-tx,
bfd-min-rx, bfd-multiplier>
    --ipv4static-route-bfd  <border-leaf-ip, destination-ip, source-ip, bfd-min-tx,
bfd-min-rx, bfd-multiplier>
```

Example

The following example creates static route BFD on tenant VRF:

```
efa tenant vrf create --name VRF1 --tenant tenant1
     --ipv6-static-route-bfd BL1-IP,1001::2,1001::1,100,200,5
     --ipv6-static-route-bfd BL2-IP,1011::2,1011::1,100,200,5
     --ipv6-static-route-bfd BL1-IP,1002::2, 1002::1
     --ipv6-static-route-bfd BL2-IP,1012::2, 1012::1
     --ipv4-static-route-bfd BL1-IP,13.0.0.1,13.0.0.9,200,300,6
     --ipv4-static-route-bfd BL2-IP,13.0.1.1,13.0.1.9,200,300,6
     --ipv4-static-route-bfd BL1-IP,13.0.0.2,13.0.0.10
     --ipv4-static-route-bfd BL2-IP,13.0.1.2,13.0.1.10
```

Configure Peer Group on Tenant BGP

You can configure a peer group on a tenant BGP.

About This Task

Follow this procedure to configure a peer group.

Procedure

1. To configure BGP peer-group on tenant BGP, run the following command:

```
efa tenant service bgp peer-group create --name <peer-group-name> --tenant <tenant-
name>
      --description <description>
          --pg-name <border-leaf-ip:pg-name> --pg-asn <border-leaf-ip:pg-name,remote-
asn>
          --pg-bfd <border-leaf-ip:pg-name,bfd-enable(true/false),interval,min-
rx,multiplier>
          --pg-next-hop-self <border-leaf-ip:pg-name,next-hop-self(true/false/always)>
          --pg-update-source-ip <border-leaf-ip:pg-name,update-source-ip>
```

   The following example creates a BGP peer group on tenant BGP:

```
efa tenant service bgp peer-group create –name ten1BgpPG1 --tenant tenant1
    --pg-name BL1-IP:pg1 --pg-asn BL1-IP:pg1,6000
    --pg-bfd BL1-IP:pg1,true,100,200,5
    --pg-next-hop-self BL1-IP:pg1,true
    --pg-update-source-ip BL1-IP:pg1,10.20.30.40
```

2. To update BGP peer-group, run the following command:

```
efa tenant service bgp peer-group update --name <peer-group-name> --tenant <tenant-
name>
    --operation <peer-group-add|peer-group-delete|peer-group-desc-update> --
description <description>
    --pg-name <border-leaf-ip:pg-name> --pg-asn <border-leaf-ip:pg-name,remote-asn>
    --pg-bfd <border-leaf-ip:pg-name,bfd-enable(true/false),interval,min-rx,multiplier>
    --pg-next-hop-self <border-leaf-ip:pg-name,next-hop-self(true/false/always)>
    --pg-update-source-ip <border-leaf-ip:pg-name,update-source-ip>
```

The following example updates a BGP peer group on tenant BGP:

```
efa tenant service bgp peer-group update --name ten1BgpPG1 --tenant tenant1
     --operation peer-group-add
     --pg-name BL1-IP:pg2 -pg-asn BL1-IP:pg2,7000
     --pg-bfd BL1-IP:pg2,true,200,300,6
     --pg-next-hop-self BL1-IP:pg2,true
     --pg-update-source-ip BL1-IP:pg2,10.20.30.41
```

Configure Static Peer on Tenant BGP

You can configure a static peer on a tenant BGP.

About This Task

Follow this procedure to configure a static peer.

Procedure

1. To configure BGP static peer, run the following commands:

```
efa tenant service bgp peer create --name <peer-name> --tenant <tenant-
name>
    --ipv4-uc-nbr <border-leaf-ip,vrf-name:ipv4-neighbor,remote-
as>
    --ipv4-uc-nbr-bfd <border-leaf-ip,vrf-name:ipv4-neighbor,bfd-enable(true/false/
always),bfd-interval,bfd-rx,bfd-mult>
    --ipv4-uc-nbr-update-source-ip <border-leaf-ip,vrf-name:ipv4-neighbor,update-
source-ip>
    --ipv4-uc-nbr-next-hop-self <border-leaf-ip,vrf-name:ipv4-neighbor,next-hop-
self(true/false/always)>
    --ipv6-uc-nbr <border-leaf-ip,vrf-name:ipv6-neighbor,remote-as>
    --ipv6-uc-nbr-bfd <border-leaf-ip,vrf-name:ipv6-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfd-mult>
    --ipv6-uc-nbr-update-source-ip <border-leaf-ip,vrf-name:ipv6-neighbor,update-
source-ip>
    --ipv6-uc-nbr-next-hop-self <border-leaf-ip,vrf-name:ipv6-neighbor,next-hop-
self(true/false/always)>
```

   The following example creates a BGP static peer on tenant BGP:

```
efa tenant service bgp peer create --name bgpservice1 --tenant
tenant1
    --ipv4-uc-nbr BL1-IP,VRF1:10.20.30.40,5000
    --ipv4-uc-nbr-bfd BL1-IP,VRF1:10.20.30.40,true,100,200,5
    --ipv4-uc-nbr-update-source-ip BL1-IP,VRF1:10.20.30.40,11.22.20.33
    --ipv4-uc-nbr-next-hop-self BL1-IP,VRF1:10.20.30.40,true
```

2. To update BGP static peer, run the following commands:

```
efa tenant service bgp peer update --name <peer-name> --tenant <tenant-name>
    --operation peer-add
    --ipv4-uc-nbr <border-leaf-ip,vrf-name:ipv4-neighbor,remote-as>
    --ipv4-uc-nbr-bfd <border-leaf-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfd-mult>
    --ipv4-uc-nbr-update-source-ip <border-leaf-ip,vrf-name:ipv4-neighbor,update-
source-ip>
    --ipv4-uc-nbr-next-hop-self <border-leaf-ip,vrf-name:ipv4-neighbor,next-hop-
self(true/false/always)>
    --ipv6-uc-nbr <border-leaf-ip,vrf-name:ipv6-neighbor,remote-as>
    --ipv6-uc-nbr-bfd <border-leaf-ip,vrf-name:ipv6-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfd-mult>
    --ipv6-uc-nbr-update-source-ip <border-leaf-ip,vrf-name:ipv6-neighbor,update-
source-ip>
    --ipv6-uc-nbr-next-hop-self <border-leaf-ip,vrf-name:ipv6-neighbor,next-hop-
self(true/false/always)>
```

The following example updates a BGP static peer on tenant BGP:

```
efa tenant service bgp peer update --name bgpservice1 --tenant tenant1 --operation
peer-add
    --ipv6-uc-nbr BL1-IP,VRF1:10::40,5000
    --ipv6-uc-nbr-bfd BL1-IP,VRF1:10::40,true,100,200,5
    --ipv6-uc-nbr-update-source-ip BL1-IP,VRF1:10::40,11::22
    --ipv6-uc-nbr-next-hop-self BL1-IP,VRF1:10::40,true
```

Configure Dynamic Peer on Tenant BGP

You can configure a dynamic peer on a tenant BGP.

About This Task

Follow this procedure to configure a dynamic peer.

Procedure

1. To configure BGP dynamic peer, run the following commands:

```
efa tenant service bgp peer create --name <peer-name> --tenant <tenant-name>
    --ipv4-uc-dyn-nbr <border-leaf-ip,vrf-name:listen-range,peer-group-name,listen-
limit>
    --ipv6-uc-dyn-nbr <border-leaf-ip,vrf-name:listen-range,peer-group-name,listen-
limit>
```

The following example creates a BGP dynamic peer on tenant BGP:

```
efa tenant service bgp peer create --name bgpservice1 --tenant tenant1
    --ipv4-uc-dyn-nbr BL1-IP,VRF1:11.22.33.44/30,pg1,10
```

2. To update BGP dynamic peer, run the following commands:

```
efa tenant service bgp peer update --name <peer-name> --tenant <tenant-name>
    --operation peer-add
    --ipv4-uc-dyn-nbr <border-leaf-ip,vrf-name:listen-range,peer-group-name,listen-
limit>
    --ipv6-uc-dyn-nbr <border-leaf-ip,vrf-name:listen-range,peer-group-name,listen-
limit>
```

The following example updates a BGP dynamic peer on tenant BGP:

```
efa tenant service bgp peer create --name bgpservice1 --tenant
tenant1
    --operation peer-add –ipv4-uc-dyn-nbr BL1-IP,VRF1:11::22/127,pg1,20
```

Centralized Routing on Single Rack Small Data Center Leaf Pair (not Border Leaf Pair)

The following items are required before you configure centralized routing on Single Rack Small Data Center Leaf Pair.

- The device-role (leaf or border-leaf) are specified during the addition of the devices to the fabric, prior to the “fabric configure”.
- Border-Leaf pair can exist in a Clos or Small Data Center fabric irrespective of the VRFs instantiated in the fabric are distributed or centralized.
- Device role border-leaf implies the leaf pair used at the edge (border) of the fabric, and not restricted to the centralized routing.
- Tenant (PO, VRF, EPG, or BGP) provisioning happens on a configured fabric.
- Only the Border-Leaf devices can act as Centralized Routers.

- Default routing-type for a VRF is “distributed” and you need to explicitly provide the value “centralized” if needed.
- During creation of VRF as a CR (Centralized Router), XCO must instantiate the VRF on a pair of Border-Leaf devices.
- If the fabric (Clos or Small Data Center) has only one pair of Border-Leaf devices, then the same pair will be chosen as the designated CRs (Centralized Routers) for the VRF. Otherwise, you must explicitly provide the Border-Leaf devices as the designated CRs (Centralized Routers) during the creation of VRF.
- XCO is designed to expand or compress with the addition or deletion of racks (rack = MCT-pair) as per your requirement.
- XCO cannot determine “a given fabric is a single rack small data center fabric and can never be expanded beyond that”. Hence there is no specific automation for a single-rack use case.
- For CR on a single rack small data center fabric, as a best practice, you must configure the fabric with the device-role = border-leaf for both the MCT nodes.
- Using the **Day 1 Centralized Routing provisioning on a “Day 0 Configured Single Rack Leaf Small Data Center Fabric”** results in failure because CR can be instantiated only on the border-leaf pair of small data center fabric.
- You cannot recreate the fabric with device-role = border-leaf.

Fabric Setting for a Single Rack Deployment

1. Use the following command to configure a single-rack-deployment when you update a fabric setting:
   ```
   efa fabric setting update --name <non-clos-fabric-name> --single-rack-deployment <Yes|No>
   ```
2. The fabric setting is applicable only for a Small Data Center fabric.
3. Default value of single-rack-delployment is No.
4. **Single Rack Deployment**
   a. When the value of single-rack-deployment is Yes and the fabric is configured,
      - You cannot modify the value of single-rack-deployment from Yes to No.
      - The state is used as an indicator to XCO that “a given fabric is a single rack non-Clos fabric and will never be expanded beyond that”, so that XCO can have specific automation for the specific scenario of allowing the non-border-leaf rack to act as CR (Centralized Router) for single rack small data center leaf pair deployments.
      - You cannot expand such a fabric. If you intend to expand such a fabric, then you must delete the fabric and recreate the same with “single-rack-deployment = No”.
   b. When the value of single-rack-deployment is Yes and the fabric is not configured,
      - You can modify the single-rack-deployment value from Yes to No.

5. Multi Rack Deployment
   a. When the value of single-rack-deployment is No and the fabric is configured or not-configured,
      - You can modify the value of single-rack-deployment from No to Yes, provided the existing number of rack in the fabric is 1.
6. Fabric Device Add
   - Validations are done to ensure the number of racks in the given fabric adhere to the fabric settings.

Create a Tenant VRF for Single Rack Small Data Center Leaf Pair Deployment

You can create a tenant VRF on a single rack small data center deployment.

Before You Begin

Everything about the VRF remains as it is except for the Centralized Routing usecase on Single Rack Small Data Center Leaf Pair Deployment.

About This Task

Follow this procedure to configure a tenant VRF on single rack small data center leaf pair deployment.

Procedure

1. Create a VRF on Multi Rack Small Data Center or Multi Stage Clos Deployment.
2. Create a VRF on Single Rack Small Data Center Border-Leaf Pair Deployment.
3. Create a VRF on Single Rack Small Data Center Leaf Pair Deployment (single-rack-deployment = No).
4. Create a VRF on Single Rack Small Data Center Leaf Pair Deployment (single-rack-deployment = Yes).
   a. Create VRF on Distributed Router.
   b. Create VRF on Centralized Router.
      - During the creation of CR (Centralized Router) VRF, the single leaf rack will be considered as the designated CRs.
      - VRF will be instantiated as the CR VRF as input by the user and not DR (Distributed Router).

Configure a Single-Rack Leaf in Day 0 and Day 1 Provisioning

You can configure a single-rack leaf on Day 0 and then configure centralized and distributed routing on Day 1.

About This Task

Follow this procedure to configure a single-rack leaf in a small data center fabric.

Tip
If any devices in a fabric are in "admin-down" state, use of the following commands in that same fabric will not add or delete devices in the fabric: **`efa fabric device add-bulk`** and **`efa fabric device remove`**.

Procedure

1. Configure a single-rack leaf in a small data center fabric on Day 0.

```
efa fabric create –name <fabric-name> --type non-clos
efa fabric device add-bulk --name <fabric-name> --rack <rack-name> --ip <ip-pair>
--username <username> --password <password>
efa fabric configure --name <fabric-name>
```

2. Configure centralized routing on Day 1.

```
efa fabric setting update –name <fabric-name> --single-rack-deployment Yes
efa tenant vrf create –name <vrf-name> --tenant <tenant-name> --routing-type
centralized
```

3. Configure distributed routing on Day 1.

```
efa tenant vrf create –name <vrf-name> --tenant <tenant-name>
```

There are no changes in the provisioning model.

BFD Timers for Router BGP BFD and Static Route BFD Sessions

- On the SLX, you can provide the Bidirectional Forwarding Detection (BFD) timer configurations per static-route, per BGP peer and per BGP peer-group.

  In SLX versions prior to 20.3.2:

  1. The timer values provided per static-route, per BGP peer, and per BGP peer-group are not effective for the single-hop static route BFD sessions and single-hop router BGP BFD sessions.
  2. For the timer values to be effective for the single-hop sessions, you must additionally configure the BFD timers also on the source interface (the interface acting as the BFD source).

  But this restriction is removed from 20.3.2 onwards.
- On the SLX, you can provide the BFD timer values directly at the “router bgp” level.

  1. In SLX versions prior to 20.3.2:

     The timer values provided directly at the “router bgp” level are effective only for the multi-hop BFD sessions of both default and non-default VRF BGP peers.
  2. In SLX versions 20.3.2 or above:

The timer values provided directly at the “router bgp” level are effective for both multi-hop and single-hop BFD sessions of both default and non-default VRF BGP peers.

Note

- Prior to upgrade of SLX-OS, ensure that you enable the maintenance mode on reboot using the **`efa inventory device setting update --maint-mode-enable-on-reboot Yes –ip <IP address of SLX>`** command. If you have upgraded SLX-OS without enabling the maintenance mode, trigger a manual drift and reconcile (DRC) using the command **`efa inventory drift-reconcile execute --ip <IP address of SLX>” –reconcile`**.
- If you have not triggered the manual DRC, the **`efa fabric show --name <fabric name>`** command shows devices in `cfg-refreshed` state, and the intended BFD configuration does not get configured on the devices.
- For the static-route and router-bgp BFD enhancement to be effective, you must upgrade SLX to 20.3.2 or later.

Configure Next Hop Recursion

You can enable next hop recursion (NHR) on each tenant VRF when you create or update a tenant VRF on the switches.

Before You Begin

- By default, the NHR is disabled.
- When you upgrade to XCO 3.2.1 or later, the NHR gets disabled on VRF.
- SLXOS 20.5.1 and later supports the next hop recursion configuration on tenant VRF. For details on hardware support, refer to the SLX-OS documentation.

About This Task

Follow this procedure to configure next hop recursion.

Based on the endpoints present in the EPG, VRF is instantiated on the switches when you create an L3 endpoint group or transition an endpoint group to L3 endpoint group. Based on the endpoints present in the EPG, VRF is updated on the switches when you update a VRF.

The next hop recursion is configured when you configure a VRF.

Procedure

1. Run the following command to enable next hop recursion (NHR) when you create a tenant VRF:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name> --next-hop-recursion-enable {true|false}
```

2. Run the following command to enable next hop recursion when you update a tenant VRF:

```
efa tenant vrf update --name <vrf-name> --tenant <tenant-name>
            --operation next-hop-recursion-update –next-hop-recursion-enable {true|false}
```

Example

```
efa tenant vrf create --name vs --tenant t1 --next-hop-recursion-enable true
./efa tenant vrf show --tenant t1 --name vs --detail
Name                        : vs
Tenant                      : t1
Routing Type                : distributed
Centralized Routers         :
Enable Layer3 Extension     : true
Redistribute                : connected
Max Path                    : 8
Local Asn                   :
L3VNI                       :
EVPN IRB BD                 :
EVPN IRB VE                 :
BR VNI                      :
BR BD                       :
BR VE                       :
RH Max Path                 :
Enable RH ECMP              : false
Enable Graceful Restart     : false
Enable NextHop Recursion    : true
Route Target                :
Static Route                :
Static Route BFD            :
Network Route Address       :
Static Network              :
Aggregate Address           :
VRF Type                    : private
State                       : vrf-created
Dev State                   : not-provisioned
App State                   : cfg-ready




efa tenant epg create --name epg1 --tenant t1 --switchport-mode trunk --po po1
--port 10.20.246.15[0/18] --vrf vs --l3-vni 30211 --ctag-range 23-25 --anycast-ip
23:23.10.12.2/24  --anycast-ip 24:24.10.12.1/24 --anycast-ip 25:25.10.12.1/24  --suppress-
arp 25:true

efa tenant vrf show --tenant t1 --name vs --detail
Name                        : vs
Tenant                      : t1
Routing Type                : distributed
Centralized Routers         :
Enable Layer3 Extension     : true
Redistribute                : connected
Max Path                    : 8
Local Asn                   :
L3VNI                       : 30211
EVPN IRB BD                 : 4096
EVPN IRB VE                 : 8192
BR VNI                      :
BR BD                       :
BR VE                       :
RH Max Path                 :
Enable RH ECMP              : false
Enable Graceful Restart     : false
Enable NextHop Recursion    : true
Route Target                : import 101:101
                            : export 101:101
Static Route                :
Static Route BFD            :
```

```
Static Network              :
Aggregate Address           :
VRF Type                    : private
State                       : vrf-device-created
Dev State                   : provisioned
App State                   : cfg-in-sync
```

```
Rack1-Device1# show run router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.3 remote-as
4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.206/32
  network 172.31.254.222/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
Rack1-Device1#
```

```
Rack1-Device2# show run router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.2 remote-as
4200000000
 neighbor 10.20.20.2 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.182/32
  network 172.31.254.222/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
Rack1-Device2#
```

```
efa tenant vrf update --name vs --tenant t1 --operation next-hop-recursion-update --next-
hop-recursion-enable false


efa tenant vrf show --tenant t1 --name vs --detail
Name                        : vs
Tenant                      : t1
Routing Type                : distributed
Centralized Routers         :
Enable Layer3 Extension     : true
Redistribute                : connected
Max Path                    : 8
Local Asn                   :
L3VNI                       : 30211
EVPN IRB BD                 : 4096
EVPN IRB VE                 : 8192
BR VNI                      :
BR BD                       :
BR VE                       :
RH Max Path                 :
Enable RH ECMP              : false
```

```
Enable Graceful Restart     : false
Enable NextHop Recursion    : false
Route Target                : import 101:101
                            : export 101:101
Static Route                :
Static Route BFD            :
Network Route Address       :
Static Network              :
Aggregate Address           :
VRF Type                    : private
State                       : vrf-device-created
Dev State                   : provisioned
App State                   : cfg-in-sync
```

```
Rack1-Device1# show run router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.3 remote-as
4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.206/32
  network 172.31.254.222/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
Rack1-Device1#
```

```
Rack1-Device2# show run router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.2 remote-as
4200000000
 neighbor 10.20.20.2 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.182/32
  network 172.31.254.222/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  next-hop-recursion
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
Rack1-Device2#
```

Provision a Tenant Endpoint Group

You can provision a tenant endpoint group.

About This Task

Complete the following tasks to provision a tenant endpoint group in your XCO fabric.

Procedure

1. Create a Tenant Endpoint Group on page 227
2. Update a Tenant Endpoint Group on page 227
3. Show a Tenant Endpoint Group on page 228
4. Delete a Tenant Endpoint Group on page 229

5. Configure Network Property Description on Tenant EPG on page 229
6. Enable or Disable ICMP Redirect on Tenant EPG Networks on page 251
7. Update Anycast IP on an Existing Tenant Network on page 268
8. Configure Multiple Anycast IP on page 271
9. Configure IPv6 Neighbor Discovery (ND) on a Tenant Network on page 274
10. Configure BFD Session Type for an Endpoint Group on page 276
11. Configure Cluster Edge Port (CEP) Cluster Tracking for Endpoint Groups on page 277
12. Configure Suppress Address Resolution Protocol and Neighbor Discovery on VLAN or Bridge Domain on page 278
13. Configure Local IP for Endpoint Group on page 282
14. IPv6 Anycast Gateway Support on page 207
15. Software BFD Session Support on CEP on page 285

Create a Tenant Endpoint Group

An endpoint group is a logical group of endpoints, which are devices that are connected to the network. You can specify parameters, such as group name, IP address, port channels, switchport mode, BGP service type, native VLAN, CTAG range, associated VRF, Layer 2 and Layer 3 VNI, bridge domain, and neighbor discovery preferences.

About This Task

Follow this procedure to create a tenant endpoint group.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To create a tenant endpoint group, run the following command:

```
efa tenant epg create --name ten1epg1 --tenant tenant1
```

Update a Tenant Endpoint Group

When you update a tenant endpoint group, you can specify the resources like device ports, VLAN range, L2 VNI range, L3 VNI range, and VRF count.

About This Task

Follow this procedure to update a tenant endpoint group.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To update a tenant endpoint group when you create a tenant, run the following command:

```
efa tenant create --name <tenant-name> --description <tenant-description> --l2-vni-range
<value> --l3-vni-range <value> -- vlan-range <value> --vrf-count <value> --enable-bd –
port <list-of-ports>
```

Example

1. The following example creates a VLAN-based tenant with manual VNI mapping:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant11 --l2-vni-range
10002-14190
--l3-vni-range 14191-14200 --vlan-range 2-4090 --vrf-count 10 --port
10.20.216.15[0/11-20],10.20.216.16[0/11-20]
--description Subscriber1

Tenant created successfully.

--- Time Elapsed: 455.141597ms ---
```

2. The following example creates a BD-based tenant:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant21 --l2-vni-range
30002-34190
--l3-vni-range 34191-34200 --vlan-range 2-4090 --vrf-count 10 --enable-bd
--port 10.20.216.15[0/21-28],10.20.216.16[0/21-28]

Tenant created successfully.

--- Time Elapsed: 501.176996ms ---
```

3. The following example creates a tenant with auto-VNI with breakout ports:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name tenant12
--vlan-range 2-100 --vrf-count 10 --port
10.20.216.103[0/1-10],10.20.216.104[0/1-5,0/6:1-4]

Tenant created successfully.

--- Time Elapsed: 427.73527ms ---
```

4. The following creates a shared tenant:

```
(efa:extreme)extreme@node-1:~$ efa tenant create --name ST
--type shared --port 10.20.216.15[0/1-10],10.20.216.16[0/1-10]

Tenant created successfully.

--- Time Elapsed: 381.182892ms ---
```

Show a Tenant Endpoint Group

An endpoint group is a logical group of endpoints, which are devices that are connected to the network. You can specify such parameters as group name, the IP address, the port channels, the switchport mode, the BGP service type, native VLAN, CTAG range, the associated VRF, the Layer 2 and Layer 3 VNI, the bridge domain, and neighbor discovery preferences.

About This Task

Follow this procedure to show a tenant endpoint group.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To show a tenant endpoint group, run the following command:

```
efa tenant epg create --name ten1epg1 --tenant tenant1
```

Example

Delete a Tenant Endpoint Group

An endpoint group is a logical group of endpoints, which are devices that are connected to the network. You can specify such parameters as group name, the IP address, the port channels, the switchport mode, the BGP service type, native VLAN, CTAG range, the associated VRF, the Layer 2 and Layer 3 VNI, the bridge domain, and neighbor discovery preferences.

About This Task

Follow this procedure to delete a tenant endpoint group.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

To delete a tenant endpoint group, run the following command:

```
efa tenant epg create --name ten1epg1 --tenant tenant1
```

Example

Configure Network Property Description on Tenant EPG

The EPG (endpoint group) Network Property Description enables you to configure “description” for each XCO tenant ctag which gets configured on the SLX as VLAN or BD description.

About This Task

Follow this procedure to configure network property description on tenant EPG network.

The following table describes the default value of “description”:

Network Type	Description
L2 Extension	Tenant L2 Extended VLAN or BD
L3 Extension	Tenant L3 Extended VLAN or BD
L3 Hand off	Tenant L3 Hand-off VLAN or BD
L3 Extension EVPN IRB	Tenant L3 Extended IRB BD

You can provide the EPG network “description” when you create or update an EPG (ctag-range-add).

Procedure

1. To configure description when you create a tenant EPG, run the following command:

```
efa tenant epg create --name <epg-name> --tenant string <tenant-name> --port <list-of-
phy> --po <list-of-po>
       --switchport-mode <access |trunk | trunk-no-default-native>
       --ctag-range string <ctag-range> --ctag-description <ctag:description>
```

2. To configure description when you update a tenant EPG, run the following command:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name> --operation <ctag-range-
add>
      --ctag-range <ctag-range> --ctag-description <ctag:description>
```

   The following example shows network property:

```
efa tenant show
+--------+--------+-------+-------+------+------+-------+----------------------+
|  Name  |  Type  | VLAN  | L2VNI | L3VNI| VRF  | Enable|         Ports        |
|        |        | Range | Range | Range| Count| BD    |                      |
+--------+--------+-------+-------+------+------+-------+----------------------+
| bdTen1 | private| 21-30 |       |      |  10  | true  | 10.20.246.15[0/11-20]|
|        |        |       |       |      |      |       | 10.20.246.16[0/11-20]|
+--------+--------+-------+-------+------+------+-------+----------------------+
|vlanTen1| private| 11-20 |       |      |  10  | false | 10.20.246.16[0/1-10] |
|        |        |       |       |      |      |       | 10.20.246.15[0/1-10] |
+--------+--------+-------+-------+------+------+-------+----------------------+
efa tenant vrf create -name ten1vrf1 -tenant vlanTen1

efa tenant vrf create -name ten2vrf2 -tenant bdTen1
```

3. Run the following show command:

```
efa tenant epg show --detail
================================================================================================
====================================
Name          : epg2
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/35]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto

NW Policy     : Ctag Range                    : 360
              : VRF                           : VRF11
              : L3Vni                         : 15191

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+--------------------+-------------+-------------+
|        Port        |  Dev State  |  App State  |
+--------------------+-------------+-------------+
| 10.20.246.15[0/35] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
Port Property States
```

```
+----+-------------+------+----+-----------+-------+---------+------------+----+----
+-------+-------+-----------+-----------+
|Ctag|    Ctag     |L2Vni |BD  |Anycast    |Anycast| Suppress|  Local IP  | IP |IPv6|
IPv6 ND|IPv6 ND| Dev State | App State |
|    | Description |      |Name|IPv4       |IPv6   |  ARP/ND |[Device-IP->| MTU|ND  |
Managed|Other  |           |           |
|    |             |      |    |           |       |         | Local-IP]  |    |MTU |
Config |Config |           |           |
+----+-------------+------+----+-----------+-------+---------+------------+----+----
+-------+-------+-----------+-----------+
|360 |Tenant L3    |11003 |    |36.1.1.1/24|       |   T/F   |            |    |    |
false | false |provisioned|cfg-in-sync|
|    |Extended VLAN|      |    |           |       |         |            |    |
|       |       |           |           |
+----+-------------+------+----+-----------+-------+---------+------------+----+----
+-------+-------+-----------+-----------+
Network Property [Flags : * - Native Vlan]

+-----+--------+----------+---------+-----------+------------+
| Ctag| IPv6 ND| No       | Valid   | Preferred | Config Type|
|     | Prefix | Advertise| Lifetime| Lifetime  |            |
+-----+--------+----------+---------+-----------+------------+
IPv6 ND Prefix Flags

+-----+---------------+------------+---------------+-----------+-----------+
|Ctag |  MAC ACL IN   |MAC ACL OUT |   IP ACL IN   |IP ACL OUT |IPv6 ACL IN|
+-----+---------------+------------+---------------+-----------+-----------+
|360  |ext-mac-permit |            |ext-ip-permit  |           |           |
|     |-any-mirror-acl|            |-any-mirror-acl|           |           |
+-----+---------------+------------+---------------+-----------+-----------+
Network Property ACLs
```

4. Verify the switch configuration on SLX device.

```
Rack1-Device1# show running-
config bridge-domain
bridge-omain 1 p2mp
 description Tenant L2 Extended BD
 pw-profile default
 logical-interface ethernet 0/11.21
 bpdu-drop-enable
 local-switching
!
bridge-domain 2 p2mp
 description Ten2BDNW1
 pw-profile default
 logical-interface ethernet 0/11.22
 bpdu-drop-enable
 local-switching
!
bridge-domain 3 p2mp
 description Tenant L3 Extended BD
 pw-profile Tenant-profile
 router-interface Ve 4099
 !
 logical-interface ethernet 0/11.23
 bpdu-drop-enable
 local-switching
 suppress-arp
!
bridge-domain 4 p2mp
 description Ten2BDNW2
 pw-profile Tenant-profile
 router-interface Ve 4100
 !
 logical-interface ethernet 0/11.24
 bpdu-drop-enable
 local-switching
 suppress-arp
!
bridge-domain 4093 p2mp
 description Tenant L3 Extended IRB BD
 pw-profile Tenant-profile
 router-interface Ve 8189
 !
 bpdu-drop-enable
 local-switching
!
bridge-domain 4094 p2mp
 description Tenant L3 Extended IRB BD
 pw-profile Tenant-profile
 router-interface Ve 8190
 !
 bpdu-drop-enable
 local-switching
!
```

```
Rack1-Device2# show running-
config bridge-domain
bridge-domain 1 p2mp
 description Tenant L2 Extended BD
 pw-profile default
 logical-interface ethernet 0/11.21
 bpdu-drop-enable
 local-switching
!
bridge-domain 2 p2mp
 description Ten2BDNW1
 pw-profile default
 logical-interface ethernet 0/11.22
 bpdu-drop-enable
 local-switching
!
bridge-domain 3 p2mp
 description Tenant L3 Extended BD
 pw-profile Tenant-profile
 router-interface Ve 4099
 !
 logical-interface ethernet 0/11.23
 bpdu-drop-enable
 local-switching
 suppress-arp
!
bridge-domain 4 p2mp
 description Ten2BDNW2
 pw-profile Tenant-profile
 router-interface Ve 4100
 !
 logical-interface ethernet 0/11.24
 bpdu-drop-enable
 local-switching
 suppress-arp
!
bridge-domain 4093 p2mp
 description Tenant L3 Extended IRB BD
 pw-profile Tenant-profile
 router-interface Ve 8189
 !
 bpdu-drop-enable
 local-switching
!
bridge-domain 4094 p2mp
 description Tenant L3 Extended IRB BD
 pw-profile Tenant-profile
 router-interface Ve 8190
 !
 bpdu-drop-enable
 local-switching
!
```

```
Rack1-Device1#show running-config vlan
vlan 11
 description Tenant L2 Extended VLAN
!
vlan 12
 description Ten1VLANNW1
!
vlan 13
```

```
Rack1-Device2# show running-config vlan
vlan 11
 description Tenant L2 Extended VLAN
!
vlan 12
 description Ten1VLANNW1
!
vlan 13
```

```
 router-interface Ve 13
 suppress-arp
 description Tenant L3 Extended VLAN
!
vlan 14
 router-interface Ve 14
 suppress-arp
 description Ten1VLANNW2
!
vlan 15
 description Tenant L3 Hand-off VLAN
!
vlan 16
 description Ten1VLANNW3
!

Rack1-Device1#
```

```
 router-interface Ve 13
 suppress-arp
 description Tenant L3 Extended VLAN
!
vlan 14
 router-interface Ve 14
 suppress-arp
 description Ten1VLANNW2
!
vlan 15
 description Tenant L3 Hand-off VLAN
!
vlan 16
 description Ten1VLANNW3
!
Rack1-Device2#
```

The following examples configure network property "description" on a tenant EPG network:

- VLAN Based L2 Extension EPG

```
efa tenant epg create --name ten1epg1 --tenant vlanTen1 --port
10.20.246.15[0/1],10.20.246.16[0/1]
--switchport-mode trunk --ctag-range 11-12 --ctag-description 12:Ten1VLANNW1
```

- VLAN Based L3 Extension EPG

```
efa tenant epg create --name ten1epg2 --tenant vlanTen1 --port
10.20.246.15[0/1],10.20.246.16[0/1]
--switchport-mode trunk --ctag-range 13-14 --ctag-description 14:Ten1VLANNW2 --
anycast-ip
13:10.0.13.1/24 --anycast-ip 14:10.0.14.1/24 --vrf ten1vrf1
```

- VLAN Based L3 Hand-off EPG

```
efa tenant epg create --name ten1epg3 --tenant vlanTen1 --type l3-hand-off
--port 10.20.246.15[0/1],10.20.246.16[0/1] --switchport-mode trunk --ctag-range
15-16 --ctag-description 16:Ten1VLANNW3
```

- BD Based L2 Extension EPG

```
efa tenant epg create --name ten2epg1 --tenant bdTen1 --port
10.20.246.15[0/11],10.20.246.16[0/11]
--switchport-mode trunk --ctag-range 21-22 --ctag-description 22:Ten2BDNW1
```

- BD Based L3 Extension EPG

```
efa tenant epg create --name ten2epg2 --tenant bdTen1 --port
10.20.246.15[0/11],10.20.246.16[0/11]
--switchport-mode trunk --ctag-range 23-24 --ctag-description 24:Ten2BDNW2 --
anycast-ip
23:10.0.23.1/24 --anycast-ip 24:10.0.24.1/24 --vrf ten2vrf
```

Configure Network Property on Tenant EPG

You can configure network properties on a tenant EPG network.

About This Task

Using the new EPG update operation `network-property-add`, `network-property-delete`, and `network-property-update` to add, delete, and update the network-property (NP) of an EPG networks. For example, If an EPG doesn't have the NP MAC

ACL applied and if you want to apply NP MAC ACL on the EPG networks, then use the `network-property-add` or `network-property-update` operation.

Note
The network property configuration on Tenant EPG is supported only for PP ACL.

Procedure

1. Run the following command to add the network property when you update an EPG network:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
   --operation network-property-add
   --switchport-native-vlan <2-4090> --l2-vni <ctag:l2-vni>
   --ip-mtu <ctag:ip-mtu> --anycast-ip <ctag:anycast-ip> --anycast-ipv6 <ctag:anycast-
ipv6
   --bridge-domain <ctag:bridge-domain> --ctag-description <ctag:vlandescription>
   --local-ip <ctag,device-ip:local-ip> --local-ipv6 <ctag,device-ip:local-ipv6>
   --ipv6-nd-mtu <ctag:mtu> --ipv6-nd-managed-config <ctag:ipv6-nd-managed-config>
   --ipv6-nd-other-config <ctag:ipv6-nd-other-config> --ipv6-nd-prefix
<ctag:prefix1,prefix2
   --ipv6-nd-prefix-valid-lifetime <ctag,prefix:validTime>
   --ipv6-nd-prefix-preferred-lifetime <ctag,prefix:preferredTime>
   --ipv6-nd-prefix-no-advertise <ctag,prefix:noadvertiseflag>
   --ipv6-nd-prefix-config-type <ctag,prefix:configType>
   --suppress-arp <ctag:suppress-arp>
   --suppress-nd <ctag:suppress-nd>
   --np-mac-acl-in <ctag:acl-name> --np-mac-acl-out <ctag:acl-name>
   --np-ip-acl-in <ctag:acl-name> --np-ip-acl-out <ctag:acl-name>
   --np-ipv6-acl-in <ctag:acl-name>
```

Example

```
efa tenant epg update --tenant t1 --name epg2 --operation network-property-add
    --np-mac-acl-in 360:ext-mac-permit-any-mirror-acl --np-ip-acl-in 360:ext-ip-permit-
any-mirror-acl

efa tenant epg show --detail
==========================================================================
Name           : epg2
Tenant         : t1
Type           : extension
State          : epg-with-port-group-and-ctag-range
Description    :

Ports          : 10.20.246.15[0/35]
POs            :
Port Property  : SwitchPort Mode               : trunk
               : Native Vlan Tagging           : false
               : Single-Homed BFD Session Type : auto

NW Policy      : Ctag Range                    : 360
               : VRF                           : VRF11
               : L3Vni                         : 15191

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+--------------------+-------------+-------------+
```

```
|        Port        |  Dev State  |  App State  |
+--------------------+-------------+-------------+
| 10.20.246.15[0/35] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
Port Property States

+-----+-------------+------+-----+------------+-------+--------+------------+------
+--------+--------------+-------------+-----------+-----------+
|Ctag |    Ctag      |L2Vni |BD   |Anycast     |Anycast|Suppress|  Local IP
|IP MTU| IPv6 ND|   IPv6 ND     |   IPv6 ND   |Dev State  |App State  |
|     | Description  |      |Name |IPv4        |IPv6   | ARP/ND |[Device-IP-
>|      |   MTU  |Managed Config |Other Config |           |           |
|     |              |      |     |            |       |        |Local-IP]
|      |        |              |             |           |           |
+-----+-------------+------+-----+------------+-------+--------+-------------------
+--------+--------------+-------------+-----------+-----------+
|360  |Tenant L3     |11003 |     |36.1.1.1/24 |       |  T/F   |
|      |        |    false      |    false    |provisioned|cfg-in-sync|
|     |Extended VLAN |      |     |            |       |        |
|      |        |              |             |           |           |
+-----+-------------+------+-----+------------+-------+--------+------------+------
+--------+--------------+-------------+-----------+-----------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid
Lifetime | Preferred Lifetime | Config Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+-----+------------------------------+--------+-----------------------------+--------
+-------+
| Ctag|         MAC ACL IN          |MAC     |
IP ACL IN          |IP      | IPv6  |
|     |                              |ACL OUT |
|ACL OUT | ACL IN|
+-----+------------------------------+--------+-----------------------------+--------
+-------+
| 360 |ext-mac-permit-any-mirror-acl|        |ext-ip-permit-any-mirror-acl|
|       |
+-----+------------------------------+--------+-----------------------------+--------
```

```
+-------+
Network Property ACLs
```

```
Rack1Device1# show run vlan 360
 vlan 360
 router-interface Ve 360
 suppress-arp
mac access-group ext-mac-permit-any-mirror-acl in
 description Tenant L3 Extended VLAN
 !

 Rack1Device1# show run int ve 360
 interface Ve 360
 vrf forwarding VRF11
 ip access-group ext-ip-permit-any-mirror-acl in
 ip anycast-address 36.1.1.1/24
 no shutdown
 !
```

```
Rack1Device1# show run vlan 360
vlan 360
 router-interface Ve 360
 suppress-arp
 mac access-group ext-mac-permit-any-mirror-acl in
 description Tenant L3 Extended VLAN
!

Rack1Device2# show run int ve 360
interface Ve 360
 vrf forwarding VRF11
 ip access-group ext-ip-permit-any-mirror-acl in
 ip anycast-address 36.1.1.1/24
 no shutdown
!
```

2. Run the following command to delete the network property:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
  --operation network-property-delete
  --switchport-native-vlan <2-4090> --l2-vni <ctag:l2-vni>
  --ip-mtu <ctag:ip-mtu> --anycast-ip <ctag:anycast-ip> --anycast-ipv6 <ctag:anycast-ipv6
  --bridge-domain <ctag:bridge-domain> --ctag-description <ctag:vlandescription>
  --local-ip <ctag,device-ip:local-ip> --local-ipv6 <ctag,device-ip:local-ipv6>
  --ipv6-nd-mtu <ctag:mtu> --ipv6-nd-managed-config <ctag:ipv6-nd-managed-config>
  --ipv6-nd-other-config <ctag:ipv6-nd-other-config> --ipv6-nd-prefix
<ctag:prefix1,prefix2
  --ipv6-nd-prefix-valid-lifetime <ctag,prefix:validTime>
  --ipv6-nd-prefix-preferred-lifetime <ctag,prefix:preferredTime>
  --ipv6-nd-prefix-no-advertise <ctag,prefix:noadvertiseflag>
  --ipv6-nd-prefix-config-type <ctag,prefix:configType>
  --suppress-arp <ctag:suppress-arp>
  --suppress-nd <ctag:suppress-nd>
  --np-mac-acl-in <ctag:acl-name> --np-mac-acl-out <ctag:acl-name>
  --np-ip-acl-in <ctag:acl-name> --np-ip-acl-out <ctag:acl-name>
  --np-ipv6-acl-in <ctag:acl-name>
```

Example

```
efa tenant epg update --tenant t1 --name epg2 --operation network-property-delete
    --np-mac-acl-in 360:ext-mac-permit-any-mirror-acl --np-ip-acl-in 360:ext-ip-permit-
any-mirror-acl

efa tenant epg show --detail
=====================================================================
Name          : epg2
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/35]
POs           :

Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
```

```
NW Policy     : Ctag Range                    : 360
              : VRF                           : VRF11
              : L3Vni                         : 15191

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+--------------------+-------------+-------------+
|        Port        |  Dev State  |  App State  |
+--------------------+-------------+-------------+
| 10.20.246.15[0/35] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
Port Property States

+----+-------------+------+----+------------+-------+--------+----------------------
+---+------+--------------+-------------+-----------+-----------+
|Ctag|     Ctag    |L2Vni |BD  |Anycast IPv4|Anycast|Suppress|      Local IP
|IP |IPv6  |   IPv6 ND    |  IPv6 ND   | Dev State | App State |
|    | Description |      |Name|            |IPv6   | ARP/ND |[Device-IP->Local-
IP] |MTU|ND MTU|Managed Config|Other Config|           |           |
+----+-------------+------+----+------------+-------+--------+----------------------
+---+------+--------------+-------------+-----------+-----------+
|360 |Tenant L3    |11003 |    |36.1.1.1/24 |       |  T/F   |
|   |      |    false     |   false    |provisioned|cfg-in-sync|
|    |Extended VLAN|      |    |            |       |        |
|   |      |              |            |           |           |
+----+-------------+------+----+------------+-------+--------+----------------------
+---+------+--------------+-------------+-----------+-----------+
Network Property [Flags : * - Native Vlan]

+------+---------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid
Lifetime | Preferred Lifetime | Config Type |
+------+---------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+--------------------+-------------+------------------+------------
+-------------+
| Ctag |     MAC ACL IN     | MAC ACL
OUT |    IP ACL IN     | IP ACL OUT | IPv6 ACL IN |
+------+--------------------+-------------+------------------+------------
+-------------+
Network Property ACLs
```

```
For 'unstable' entities, run 'efa tenant po/vrf show' for detail
```

```
Rack1Device1# show run vlan 360
vlan 360
 router-interface Ve 360
 suppress-arp
 description Tenant L3 Extended
VLAN
!

Rack1Device1# show run int ve 360
interface Ve 360
 vrf forwarding VRF11
 ip anycast-address 36.1.1.1/24
 no shutdown
!
```

```
Rack1Device2# show run vlan 360
vlan 360
 router-interface Ve 360
 suppress-arp
 description Tenant L3 Extended
VLAN
!

Rack1Device2# show run int ve 360
interface Ve 360
 vrf forwarding VRF11
 ip anycast-address 36.1.1.1/24
 no shutdown
!
```

3. Run the following command to update the network property:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
   --operation network-property-update
   --switchport-native-vlan <2-4090> --l2-vni <ctag:l2-vni>
   --ip-mtu <ctag:ip-mtu> --anycast-ip <ctag:anycast-ip> --anycast-ipv6 <ctag:anycast-
ipv6
   --bridge-domain <ctag:bridge-domain> --ctag-description <ctag:vlandescription>
   --local-ip <ctag,device-ip:local-ip> --local-ipv6 <ctag,device-ip:local-ipv6>
   --ipv6-nd-mtu <ctag:mtu> --ipv6-nd-managed-config <ctag:ipv6-nd-managed-config>
   --ipv6-nd-other-config <ctag:ipv6-nd-other-config> --ipv6-nd-prefix
<ctag:prefix1,prefix2
   --ipv6-nd-prefix-valid-lifetime <ctag,prefix:validTime>
   --ipv6-nd-prefix-preferred-lifetime <ctag,prefix:preferredTime>
   --ipv6-nd-prefix-no-advertise <ctag,prefix:noadvertiseflag>
   --ipv6-nd-prefix-config-type <ctag,prefix:configType>
   --suppress-arp <ctag:suppress-arp>
   --suppress-nd <ctag:suppress-nd>
   --np-mac-acl-in <ctag:acl-name> --np-mac-acl-out <ctag:acl-name>
   --np-ip-acl-in <ctag:acl-name> --np-ip-acl-out <ctag:acl-name>
   --np-ipv6-acl-in <ctag:acl-name>
```

Example

```
efa tenant epg update --tenant t1 --name epg2 --operation network-property-update
    --np-ip-acl-out 360:ext-ip-permit-any-mirror-acl --np-ipv6-acl-in 360:ext-ipv6-
permit-any-mirror-acl


efa tenant epg show --detail
========================================================================
Name          : epg2
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/35]
POs           :

Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto

NW Policy     : Ctag Range                    : 360
              : VRF                           : VRF11
```

```
              : L3Vni                          : 15191

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+--------------------+-------------+-------------+
|        Port        |  Dev State  |  App State  |
+--------------------+-------------+-------------+
| 10.20.246.15[0/35] | provisioned | cfg-in-sync |
+--------------------+-------------+----- --------+
Port Property States

+-----+-------------+------+--------+-----------+-------------+--------
+----------------------+----+-------+-------------+-------------+-----------
+------------+
|Ctag |    Ctag      |L2Vni |BD Name |Anycast    |Anycast IPv6 |Suppress|
Local IP        |IP  |IPv6 ND|  IPv6 ND     |   IPv6 ND   | Dev State | App State  |
|     | Description  |      |Name    |IPv4       |IPv6         | ARP/ND |[Device-
IP->Local-IP] |MTU |ND MTU |Managed Config| Other Config|           |            |
+-----+-------------+------+--------+-----------+-------------+--------
+----------------------+----+-------+-------------+-------------+-----------
+------------+
|360  |Tenant L3     |11003 |        |36.1.1.1/24|
|  T/F   |                      |    |       |
false     |    false    |provisioned| cfg-in-sync|
|     |Extended VLAN |      |        |
|             |        |                      |    |
|              |             |           |            |
+-----+-------------+------+--------+-----------+-------------+--------
+----------------------+----+-------+-------------+-------------+-----------
+------------+
Network Property [Flags : * - Native Vlan]

+------+---------------+-------------+---------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No
Advertise | Valid Lifetime | Preferred Lifetime | Config Type |
+------+---------------+-------------+---------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+-----+-------+--------+-------+-------------------------------
+--------------------------------+
|Ctag |MAC    |MAC     | IP    |
IP ACL OUT          |         IPv6 ACL IN           |
|     |ACL IN |ACL OUT | ACL
IN|                               |                                |
+-----+-------+--------+-------+-------------------------------
+--------------------------------+
|360  |       |        |       |
ext-ip-permit-any-mirror-acl |ext-ipv6-permit-any-mirror-acl |
+-----+-------+--------+-------+-------------------------------
+--------------------------------+
Network Property ACLs
```

```
For 'unstable' entities, run 'efa tenant po/vrf show' for details
```

```
Rack1Device1# show run vlan 360
vlan 360
 router-interface Ve 360
 suppress-arp
 description Tenant L3 Extended
VLAN
!

Rack1Device1# show run int ve 360
interface Ve 360
 vrf forwarding VRF11
 ip access-group ext-ip-permit-any-
mirror-acl out
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 ip anycast-address 36.1.1.1/24
 no shutdown
!
```

```
Rack1Device2# show run vlan 360
vlan 360
 router-interface Ve 360
 suppress-arp
 description Tenant L3 Extended
VLAN
!

Rack1Device2# show run int ve 360
interface Ve 360
 vrf forwarding VRF11
 ip access-group ext-ip-permit-any-
mirror-acl out
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 ip anycast-address 36.1.1.1/24
 no shutdown
!
```

IP DHCP Relay on Tenant EPG

You can configure DHCP Relay Server and Gateway Configuration per tenant network.

IP DHCP relay agents are used to forward requests and responses between the DHCP clients and the DHCP servers when they are not on the same physical subnet.

IP DHCP relay agents trap the DHCP messages between the DHCP clients and the DHCP servers and generate new forwarding message.

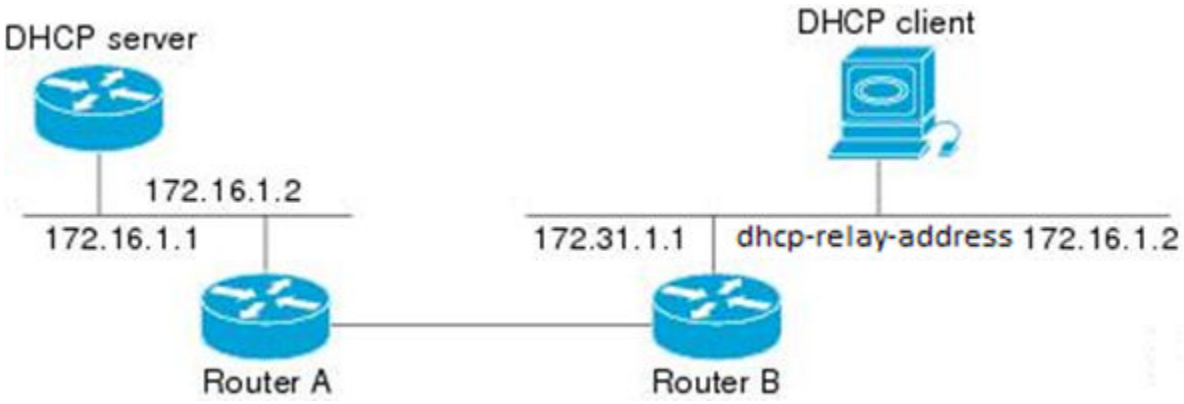

Figure 19: DHCP Client, DHCP Server, DHCP Relay

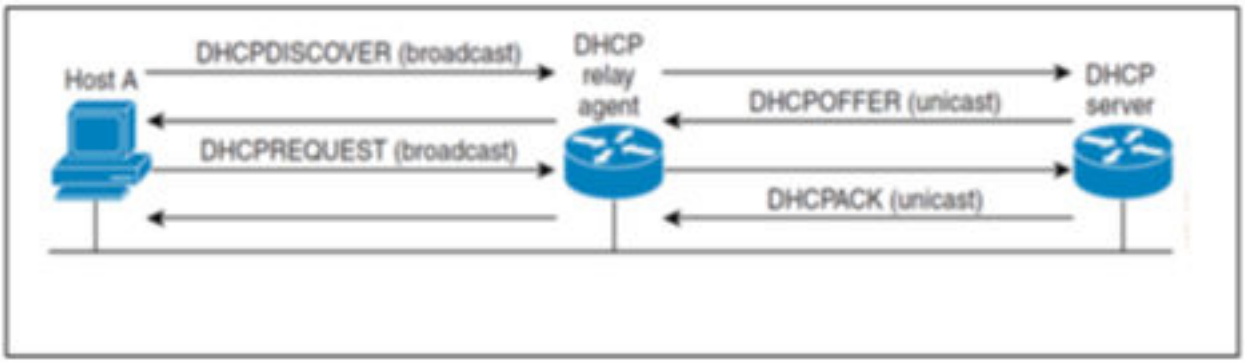

Figure 20: DHCP Relay Messages

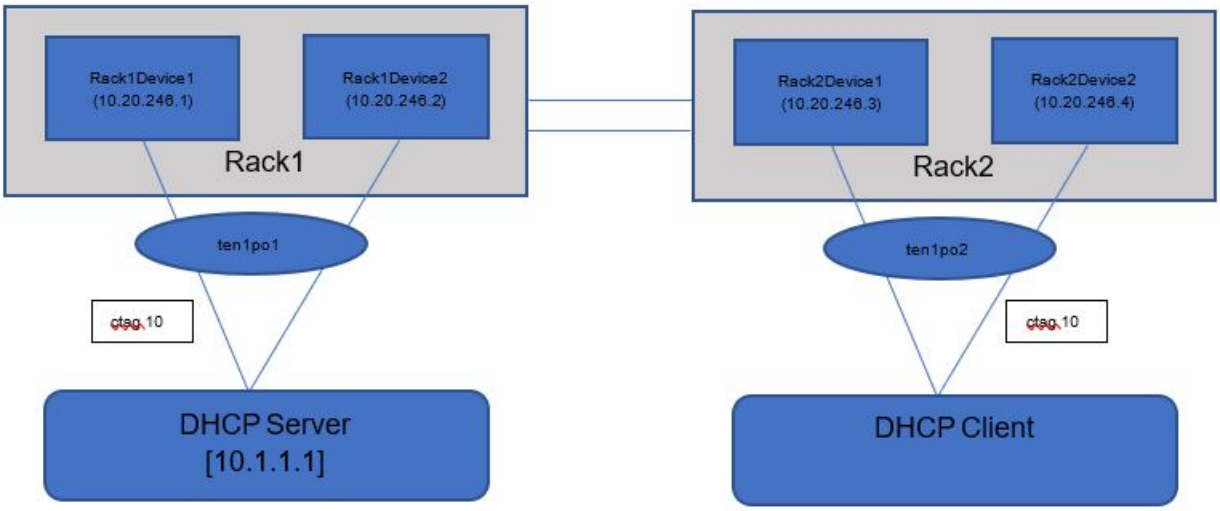

Figure 21: MultiRack Topology with DHCP Server, Client and Relay

XCO Provisioning of DHCP Relay Server and Gateway

You can provide DHCP Relay configurations (`dhcp-relay- address-ip-add/delete` and `dhcp-relay-gateway-ip-add/delete`) when you create or update an EPG.

CLI Options for DHCP Relay Server and Gateway

The following table describes the available CLI options for DHCP relay server and DHCP relay gateway (`dhcp-relay-address-ip-add/delete` and `dhcp-relay-gateway-ip-add/delete`) configurations.

DHCP Relay Server Configuration	--dhcpv4-relay-address-ip	DHCP Server IPv4 Address
	--dhcpv4-relay-address-ip-vrf	DHCP Server IPv4 Address VRF
	--dhcpv6-relay-address-ip	DHCP Server IPv6 Address
	--dhcpv6-relay-address-ip-vrf	DHCP Server IPv6 Address VRF
	--dhcpv6-relay-address-ip-interface	DHCP Server IPv6 Address Interface (eth and po)

DHCP Relay Gateway Configuration	--dhcpv4-relay-gateway-ip	DHCP Gateway IPv4 Address
	--dhcpv4-relay-gateway-ip-interface	DHCP Gateway IPv4 Address Interface **(eth)**
	--dhcpv4-relay-gateway-interface	DHCP Gateway IPv4 Interface **(eth)**
	--dhcpv6-relay-gateway-interface	DHCP Gateway IPv6 Interface **(eth)**
	--dhcpv6-relay-gateway-interface-ip	DHCP Gateway IPv6 Interface **(eth)** Address

DHCP Client and DHCP Server Residing in Same VRF

When a DHCP client and a DHCP server reside in the same VRF, use the **`dhcpv4-relay-address-ip`** and **`dhcpv6-relay-address-ip`** CLI options to provide the DHCP server address for a given tenant ctag. The DHCP server VRF is auto-derived as a VRF to which the tenant ctag belongs to.

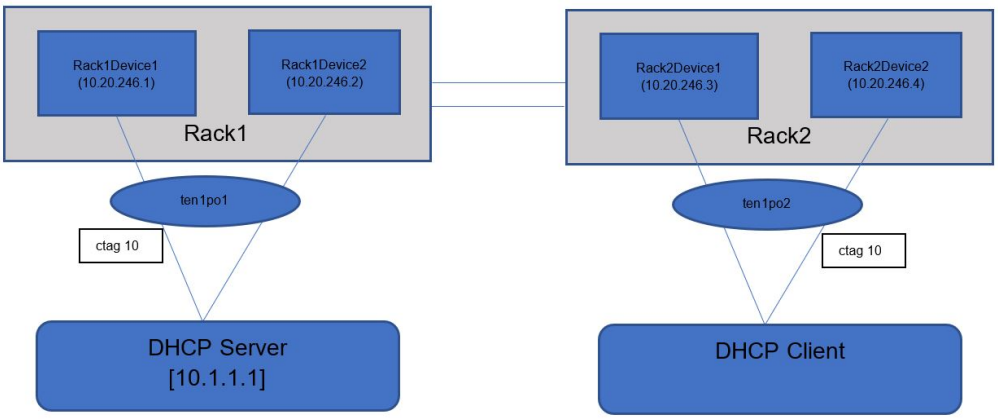

Figure 22: MultiRack Non-Clos Fabric with DHCP Server, Client, and Relay

```
efa tenant epg create --name epg1 --tenant ten1 --po
ten1po1,ten1po2
                --switchport-mode trunk --vrf vrf10 --ctag-range 10
                --anycast-ip 10:10.10.10.10/24
                --dhcpv4-relay-address-ip  10,10.20.246.1:10.1.1.1
                --dhcpv4-relay-address-ip  10,10.20.246.2:10.1.1.1

efa tenant epg show --tenant ten1  --detail --name epg1
================================================================================================
======================================
Name           : epg1
Tenant         : ten1
Type           : extension
State          : epg-with-port-group-and-ctag-range
Description    :

Ports          :
POs            : po1, po2

Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
```

```
NW Policy     : Ctag Range                    : 10
              : VRF                           : vrf10
              : L3Vni                         : 8190
+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+---------+-------------+-------------+
| Port    |  Dev State  |  App State  |
+---------+-------------+-------------+
| ten1po1 | provisioned | cfg-in-sync |
+---------+-------------+-------------+
| ten1po2 | provisioned | cfg-in-sync |
+---------+-------------+-------------+
Port Property States

+----+-------------+------+----+--------+-------+--------+----------------------+---
+------+--------------+------------+-----------+-----------+
|Ctag|     Ctag    |L2Vni |BD  |Anycast |Anycast|Suppress|      Local IP        |IP |
IPv6  |   IPv6 ND    |  IPv6 ND   | Dev State | App State |
|    | Description |      |Name|IPv4    |IPv6   | ARP/ND |[Device-IP->Local-IP] |MTU|ND
MTU|Managed Config|Other Config|           |           |
+----+-------------+------+----+--------+-------+--------+----------------------+---
+------+--------------+------------+-----------+-----------+
|360 |Tenant L3    |11003 |    |10.10.  |       |  T/F   |                      |
|      |    false     |   false    |provisioned|cfg-in-sync|
|    |Extended VLAN|      |    |10.10/24|       |        |                      |
|      |             |            |           |           |
+----+-------------+------+----+--------+-------+--------+----------------------+---
+------+--------------+------------+-----------+-----------+

Network Property [Flags : * - Native Vlan]
+------+--------------------------------+
| Ctag |           AddressIP            |
|      | Device-IP->[{Address-IP,Vrf}] |
|      |                                |
+------+--------------------------------+
|  10  |  10.20.246.1->[{10.1.1.1,}]   |
|      |  10.20.246.2->[{10.1.1.1,}]   |
+------+--------------------------------+
DHCP Relay Ips
```

```
Rack1Device1# show running-config
interface Ve
interface Ve 10
vrf forwarding vrf10
ip anycast-address 10.10.10.10/24
ip dhcp relay address 10.1.1.1
no shutdown
```

```
Rack2Device1# show running-config
interface Ve
interface Ve 10
vrf forwarding vrf10
ip anycast-address 10.10.10.10/24
no shutdown
!
```

```
Rack1Device2# show running-config
interface Ve
interface Ve 10
vrf forwarding vrf10
ip anycast-address 10.10.10.10/24
ip dhcp relay address 10.1.1.1
no shutdown
!
```

```
Rack2Device2# show running-config
interface Ve
interface Ve 10
vrf forwarding vrf10
ip anycast-address 10.10.10.10/24
no shutdown
!
```

DHCP Client and DHCP Server Residing in Different VRF

Use the `use-vrf` option on SLX to support the DHCP client and DHCP server when they are in different VRF.

Use the following CLI options to configure a VRF for a DHCP server address:

- **dhcpv4-relay-address-ip-vrf**
- **dhcpv6-relay-address-ip-vrf**

Note
Configure VRF route-leaking out of band (without XCO) on the switching or routing hardware, so that the DHCP Client and DHCP Server residing in different VRFs can communicate with each other.

DHCP Relay Server IPv4 and IPv6 Support

DHCP Relay Gateway supports IPv4 and IPv6.

Use the following CLI options to configure a DHCP IPv6 server address:

- **dhcpv6-relay-address-ip**
- **dhcpv6-relay-address-ip-vrf**
- **dhcpv6-relay-address-ip-interface**

Note
Ensure that the interface provided in the `dhcpv6-relay-address-ip-interface` option and the tenant ctag must belong to the same VRF.

DHCP Relay Gateway IPv4 and IPv6 Support

DHCP Relay Gateway supports IPv4 and IPv6. You can configure only one DHCP Relay Gateway per tenant network (ctag). Provide the DHCP Relay Gateway configuration when you create or update an DHCP Relay Gateway (`dhcp-relay-gateway-ip-add/delete`).

Use the following CLI options to configure a DHCP Relay Gateway:

- **dhcpv4-relay-gateway-ip**
- **dhcpv4-relay-gateway-ip-interface**
- **dhcpv4-relay-gateway-interface**
- **dhcpv6-relay-gateway-interface**
- **dhcpv6-relay-gateway-interface-ip**

Configure Port Property on Tenant EPG

You can configure port properties on a tenant EPG network.

About This Task

Use the EPG update operations `port-property-add`, `port-property-delete`, and `port-property-update` to add, delete, and update the port property (PP) of an EPG. For example, If an EPG doesn't have the PP MAC ACL applied and if you want to apply

PP MAC ACL on the EPG, then use the `port-property-add` or `port-property-update` operation.

Note
The port property configuration on Tenant EPG is supported only for PP ACL.

Procedure

1. Pre Configuration: Run the following command:

```
/GoDCApp/GoCommon/src/efa-client# efa tenant epg show --detail
================================================================================
Name           : epg1
Tenant         : t1
Type           : extension
State          : epg-with-port-group-and-ctag-range
Description    :
Ports          : 10.20.246.15[0/37]
               : 10.20.246.16[0/37]
POs            :
Port Property  : SwitchPort Mode               : trunk
               : Native Vlan Tagging           : false
               : Single-Homed BFD Session Type : auto
NW Policy      : Ctag Range                    : 300

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+--------------------+-------------+---------------+
|        Port        |  Dev State  |   App State   |
+--------------------+-------------+---------------+
| 10.20.246.15[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
| 10.20.246.16[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
Port Property States

+-----+-------------+------+-----+-------+--------+--------+----------------------+----
+-------+--------------+------------+-----------+-----------+
|Ctag |     Ctag    |L2Vni |BD   |Anycast| Anycast|Suppress|       Local IP       | IP
| IPv6  |   IPv6 ND    | IPv 6 ND   | Dev State | App State |
|     |  Description|      |Name |IPv4   | IPV6   | ARP/ND | [Device-IP->Local-IP]|
MTU| ND MTU|Managed Config|Other Config|           |           |
+-----+-------------+------+-----+-------+--------+--------+----------------------+----
+-------+--------------+------------+-----------+-----------+
|300  |Tenant L2    |11002 |     |       |        |  F/F   |                      |
|       |    false     |   false    |provisioned|cfg-in-sync|
|     |Extended VLAN|      |     |       |        |        |                      |
|       |              |            |           |           |
+-----+-------------+------+-----+-------+--------+--------+----------------------+----
+-------+--------------+------------+-----------+-----------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags
```

```
+------+-----------+------------+----------+-----------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+-----------+------------+----------+-----------+-------------+
Network Property ACLs

For 'unstable' entities, run 'efa tenant po/vrf show' for details
================================================================================
================================================================
```

Rack1Device1# show run int eth 0/37 interface Ethernet 0/37 cluster-track switchport switchport mode trunk switchport trunk allowed vlan add 300 no switchport trunk tag native-vlan no shutdown !	Rack1Device2# show run int eth 0/37 interface Ethernet 0/37 cluster-track switchport switchport mode trunk switchport trunk allowed vlan add 300 no switchport trunk tag native-vlan no shutdown !

2. Run the following command to add a port property when you update an EPG network:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
   --operation port-property-add
   --switchport-mode {access |trunk | trunk-no-default-native} --switchport-native-
vlan-tagging
   --single-homed-bfd-session-type {auto | hardware | software}
   --pp-mac-acl-in <acl-name> --pp-mac-acl-out  <acl-name>
   --pp-ip-acl-in <acl-name> --pp-ip-acl-out <acl-name>
   --pp-ipv6-acl-in  <acl-name>
```

Example

```
efa tenant epg update --tenant t1 --name epg1 --operation port-property-add
    --pp-mac-acl-in ext-mac-permit-any-mirror-acl --pp-ip-acl-in ext-ip-permit-any-
mirror-acl


efa tenant epg show --detail
=========================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/37]
              : 10.20.246.16[0/37]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false

              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 300

+------------------------------+-----------+----------------------------+-----------
+------------+
|         MAC ACL IN           |MAC ACL OUT |         IP ACL IN          | IP ACL OUT|
IPv6 ACL IN|
+------------------------------+-----------+----------------------------+-----------
+------------+
```

```
|ext-mac-permit-any-mirror-acl |            |ext-ip-permit-any-mirror-acl|
|            |
+------------------------------+------------+----------------------------+-----------
+------------+
Port Property ACLs

+--------------------+-------------+---------------+
|        Port        |  Dev State  |   App State   |
+--------------------+-------------+---------------+
| 10.20.246.15[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
| 10.20.246.16[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
Port Property States

+-----+--------------+------+----+-------+-------+---------+-----------+----+-----
+-------+--------+-----------+------------+
|Ctag |     Ctag     |L2Vni |BD  |Anycast|Anycast| Suppress|Local IP   |IP  |IPv6 |
IPv6 ND| IPv6 ND| Dev State | App State  |
|     | Description  |      |Name|IPv4   |IPV6   |  ARP/ND |[Device-IP-|MTU |ND   |
Managed| Other  |           |            |
|     |              |      |    |       |       |         |Local-IP]  |    |MTU  |
Config | Config |           |            |
+-----+--------------+------+----+-------+-------+---------+-----------+----+-----
+-------+--------+-----------+------------+
|300  |Tenant L2     |11002 |    |       |       |   F/F   |           |    |     |
false |false   |provisioned| cfg-in-sync|
|     |Extgended VLAN|      |    |       |       |         |           |    |
|       |        |           |            |
+-----+--------------+------+----+-------+-------+---------+-----------+----+-----
+-------+--------+-----------+------------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

For 'unstable' entities, run 'efa tenant po/vrf show' for details
```

```
================================================================================================
=========================================================
```

Rack1Device1# show run int eth 0/37	Rack1Device2# show run int eth 0/37
interface Ethernet 0/37 cluster-track switchport switchport mode trunk switchport trunk allowed vlan add 300 no switchport trunk tag native-vlan **mac access-group ext-mac-permit-any mirror-acl in** **ip access-group ext-ip-permit-any-mirror-acl in** no shutdown !	interface Ethernet 0/37 cluster-track switchport switchport mode trunk switchport trunk allowed vlan add 300 no switchport trunk tag native-vlan **mac access-group ext-mac-permit-any mirror-acl in** **ip access-group ext-ip-permit-any-mirror-acl in** no shutdown !

3. Run the following command to delete a port property:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
    --operation port-property-delete
    --switchport-mode {access |trunk | trunk-no-default-native} --switchport-native-vlan-tagging
    --single-homed-bfd-session-type {auto | hardware | software}
    --pp-mac-acl-in <acl-name> --pp-mac-acl-out  <acl-name>
    --pp-ip-acl-in <acl-name> --pp-ip-acl-out <acl-name>
    --pp-ipv6-acl-in  <acl-name>
```

Example

```
efa tenant epg update --tenant t1 --name epg1 --operation port-property-delete
    --pp-mac-acl-in --pp-ip-acl-in

efa tenant epg show --detail
=========================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/37]
              : 10.20.246.16[0/37]
POs           :

Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 300

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+--------------------+-------------+---------------+
|        Port        |  Dev State  |   App State   |
+--------------------+-------------+---------------+
| 10.20.246.15[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
| 10.20.246.16[0/37] | provisioned | cfg-in-sync   |
```

```
+--------------------+-------------+---------------+
Port Property States

+----+--------------+-----+-----+-------+-------+---------+------------+---+------
+--------+-------+-----------+-----------+
|Ctag|    Ctag      |L2Vni| BD  |Anycast|Anycast|Suppress |Local IP    |IP |IPv6  |
IPv6 ND |IPv6 ND| Dev State |App State  |
|    | Description  |     | Name|IPv4   |IPv6   | ARP/ND  |[Device-IP->|MTU|ND MTU|
Managed |Other  |           |           |
|    |              |     |     |       |       |         |Local-IP]   |   |      |
Config  |Config |           |           |
+----+--------------+-----+-----+-------+-------+---------+------------+---+------
+--------+-------+-----------+-----------+
|300 |Tenant L2     |11002|     |       |       |  F/F    |            |   |      |
false   |false  |provisioned|cfg-in-sync|
|    |Extended VLAN |     |     |       |       |         |            |   |
|        |       |           |           |
+----+--------------+-----+-----+-------+-------+---------+------------+---+------
+--------+-------+-----------+-----------+
Network Property [Flags : * - Native Vlan]
+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs
For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
=====
```

```
Rack1Device1# show run int eth
0/37
 interface Ethernet 0/37
 cluster-track
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
300
 no switchport trunk tag native-
vlan
 no shutdown
!
```

```
Rack1Device2# show run int eth
0/37
 interface Ethernet 0/37
 cluster-track
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
300
 no switchport trunk tag native-
vlan
 no shutdown
!
```

4. Run the following command to update a port property:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
    --operation port-property-update
    --switchport-mode {access |trunk | trunk-no-default-native} --switchport-native-
vlan-tagging
    --single-homed-bfd-session-type {auto | hardware | software}
    --pp-mac-acl-in <acl-name> --pp-mac-acl-out  <acl-name>
    --pp-ip-acl-in <acl-name> --pp-ip-acl-out <acl-name>
    --pp-ipv6-acl-in  <acl-name>
```

Example

```
efa tenant epg update --tenant t1 --name epg1 --operation port-property-update
       --pp-ip-acl-out ext-ip-permit-any-mirror-acl --pp-ipv6-acl-in ext-ipv6-permit-
any-mirror-acl


efa tenant epg show -detail
=====================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/37]
              : 10.20.246.16[0/37]
POs           :

Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto

NW Policy     : Ctag Range                    : 300

+------------+-------------+-----------+----------------------------
+---------------------------------+
| MAC ACL IN | MAC ACL OUT |
IP ACL IN |          IP ACL OUT         |          IPv6 ACL IN           |
+------------+-------------+-----------+----------------------------
+---------------------------------+
|            |             |           |ext-ip-permit-any-mirror-
acl | ext-ipv6-permit-any-mirror-acl |
+------------+-------------+-----------+----------------------------
+---------------------------------+
Port Property ACLs

+--------------------+-------------+---------------+
|        Port        |  Dev State  |   App State   |
+--------------------+-------------+---------------+
| 10.20.246.15[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
| 10.20.246.16[0/37] | provisioned | cfg-in-sync   |
+--------------------+-------------+---------------+
Port Property States

+-----+-------------+------+----+-------+-------+---------+-----------+---+----
+--------+--------+-----------+-----------+
|Ctag |    Ctag      |L2Vni |BD  |Anycast|
Anycast| Suppress|Local IP    |IP |IPv6|IPv6 ND | IPv6 ND|Dev State  |App State  |
|     | Description  |      |Name|IPv4   |
IPv6   |  ARP/ND |[Device-IP->|MTU|ND  |Managed | Other  |           |           |
|     |              |      |    |       |
|         |Local-IP]   |   |MTU |Config  | Config |           |           |
+-----+-------------+------+----+-------+-------+---------+-----------+---+----
+--------+--------+-----------+-----------+
|300  |Tenant L2     |11002 |    |
|       |   F/F   |            |   |    | false  |false   |provisioned|cfg-in-sync|
|     |Extgended VLAN|      |    |       |
|         |            |   |    |       |        |          |           |
+-----+-------------+------+----+-------+-------+---------+-----------+---+----
+--------+--------+-----------+-----------+
Network Property [Flags : * - Native Vlan]
```

```
+------+----------------+-------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix |
No Advertise | Valid Lifetime | Preferred Lifetime | Config Type |
+------+----------------+-------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+------------+-----------+-----------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+------------+-----------+-----------+-------------+
Network Property ACLs

For 'unstable' entities, run 'efa tenant po/vrf show' for details
```

```
Rack1Device1# show run int eth
0/37
interface Ethernet 0/37
 cluster-track
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
300
 no switchport trunk tag native-
vlan
 ip access-group ext-ip-permit-any-
mirror-acl out
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 no shutdown
!
```

```
Rack1Device2# show run int eth
0/37
interface Ethernet 0/37
 cluster-track
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
300
 no switchport trunk tag native-
vlan
 ip access-group ext-ip-permit-any-
mirror-acl out
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 no shutdown
!
```

Enable or Disable ICMP Redirect on Tenant EPG Networks

You can configure ICMP Redirect on tenant EPG network.

About This Task

Follow this procedure to enable or disable ICMP Redirect on tenant EPG networks.

You can enable or disable ICMP Redirect when you create or update an EPG using the port-group-add, ctag-range-add, vrf-add, and network-property-add or update operations.

Note

- XGS-based platforms (Extreme 8720, 8520, SLX 9150 and 9250) and J2-based SLX 9740 platform do not support ASIC for the ICMP Redirect.
- SLX-OS 20.5.1 does not support **IP ICMP Redirect**.
- Only DNX-based platforms (SLX 9540 and 9640) support ASIC for the **IP ICMP Redirect**.
- If you configure **IP ICMP Redirect** on supported platforms and later upgrade SLX to non-supporting platforms, then ensure to clean up the stale ICMP configuration on the existing VEs of XCO.

Procedure

1. To configure ICMP Redirect when you create an EPG, run the following command:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>
      --ip-icmp-redirect <ctag:ip-icmp-redirect>
      --ipv6-icmp-redirect <ctag:ipv6-icmp-redirect>
```

2. To configure ICMP Redirect when you update an EPG, run the following command:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
      --ip-icmp-redirect <ctag:ip-icmp-redirect>
      --ipv6-icmp-redirect <ctag:ipv6-icmp-redirect>
```

3. Verify the following configuration on SLX device.

```
Rack1-Device1# show running-config
interface Ve 19
interface Ve 19
vrf forwarding vrf1
ip anycast-address 3.33.3.3/24
ip mtu 1600
ip icmp redirect
ipv6 anycast-address 500::10/31
ipv6 address 700::10/31
ipv6 icmpv6 redirect
no shutdown
!
```

```
Rack1-Device2#show running-config
interface Ve 19
interface Ve 19
vrf forwarding vrf1
ip anycast-address 3.33.3.3/24
ip mtu 1600
ip icmp redirect
ipv6 anycast-address 500::10/31
ipv6 address 700::10/31
ipv6 icmpv6 redirect
no shutdown
!
```

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Example

The following example configures an ICMP Redirect on a tenant EPG network:

```
efa tenant create --name t1  --vrf-count 10 --vlan-range 10-40 --port
10.20.246.1[0/16-30],10.20.246.2[0/16-30] --l2-vni-range 1-40 --l3-vni-range 5000-50100
efa tenant po create --name po1 --tenant t1 --port 10.20.246.1[0/25],10.20.246.2[0/25] --
speed 1Gbps --negotiation active
efa tenant vrf create --name vrf1 --tenant t1
efa tenant epg create --tenant "t1" --name "epg1" --type extension --switchport-mode
trunk --single-homed-bfd-session-type auto --po po1 --vrf vrf1 --ctag-range 19 --l3-vni
5001 --anycast-ip 19:3.33.3.3/24 --bridge-domain 19:Auto-BD-2 --ctag-description
"19:Tenant L3 Extended BD" --l2-vni 19:2 --ip-mtu 19:1600 --ip-icmp-redirect 19:true --
ipv6-icmp-redirect 19:true


efa tenant epg update --tenant "t1" --name "epg1" --operation vrf-delete --vrf vrf1
efa tenant epg update --tenant "t1" --name "epg1" --operation vrf-add --vrf vrf1 --ctag-
range 19 --l3-vni 5001 --anycast-ip 19:3.33.3.3/24  --ctag-description "19:Tenant L3
Extended BD" --l2-vni 19:2 --ip-mtu 19:1600 --ip-icmp-redirect 19:true --ipv6-icmp-
redirect 19:true --anycast-ipv6 19:500::10/31 --local-ipv6 19,10.20.246.1:700::10/31 --
local-ipv6 19,10.20.246.2:700::10/31
efa tenant epg update --tenant "t1" --name "epg1" --operation ctag-range-add --ctag-range
20 --anycast-ip 20:4.33.3.3/24  --ctag-description "20:Tenant L3 Extended BD" --l2-vni
20:3  --ip-icmp-redirect 20:true
efa tenant epg update --tenant "t1" --name "epg1" --operation network-property-add --ipv6-
icmp-redirect 20:true --anycast-ipv6 20:600::10/31 --local-ipv6 20,10.20.246.1:800::10/31
--local-ipv6 20,10.20.246.2:800::10/31

========================= EPG CREATE==================================

efa tenant create --name t1  --vrf-count 10 --vlan-range 10-40 --port
10.20.246.1[0/16-30],10.20.246.2[0/16-30] --l2-vni-range 1-40 --l3-vni-range 5000-50100
```

```
Tenant created successfully.

--- Time Elapsed: 76.613817ms ---
efa tenant po create --name po1 --tenant t1 --port 10.20.246.1[0/25],10.20.246.2[0/25] --
speed 1Gbps --negotiation active

Port Channel created successfully.

--- Time Elapsed: 9.631186916s ---
efa tenant vrf create --name vrf1 --tenant t1

Vrf created successfully.

--- Time Elapsed: 105.271133ms ---
abc@abc-virtual-machine:~/GoDCApp/GoCommon/bin$
efa tenant show
+------+---------+------------+-------------+-------------+-----------+-----------
+-----------------------+--------------------------+
| Name |  Type   | VLAN Range | L2VNI Range | L3VNI Range | VRF Count | Enable BD
|         Ports         | Mirror Destination Ports |
+------+---------+------------+-------------+-------------+-----------+-----------
+-----------------------+--------------------------+
|  t1  | private |   10-40    |    1-40     | 5000-50100  |    10     |   false   |
10.20.246.1[0/16-30] |                          |
|      |         |            |             |             |           |           |
10.20.246.2[0/16-30] |                          |
+------+---------+------------+-------------+-------------+-----------+-----------
+-----------------------+--------------------------+
Tenant Details

--- Time Elapsed: 32.431956ms ---
efa tenant po show
+------+--------+----+-------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| Name | Tenant | ID | Speed | MTU | Negotiation | Min Link |  Lacp   |
Ports        |   State    |  Dev State  |  App State  |
|      |        |    |       |     |             |  Count   | Timeout
|                    |            |             |             |
+------+--------+----+-------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| po1  |   t1   | 1  | 1Gbps |     |   active    |    1     |  long   | 10.20.246.1[0/25]
| po-created | provisioned | cfg-in-sync |
|      |        |    |       |     |             |          |         | 10.20.246.2[0/25]
|            |            |             |
+------+--------+----+-------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
Port Channel Details

--- Time Elapsed: 58.50989ms ---
efa tenant vrf show
+------+--------+--------------+--------------------+--------------+--------------
+----------+----------+----------+-------------+-----------------+-----------+
| Name | Tenant | Routing Type | Centralized Routers |  Enable L3   | Redistribute | Max
Path | Local Asn | Enable GR |    State    |    Dev State    | App State |
|      |        |              |                     |  Extension   |
|          |          |          |             |                 |           |
+------+--------+--------------+--------------------+--------------+--------------
+----------+----------+----------+-------------+-----------------+-----------+
| vrf1 |   t1   | distributed  |                    |     true     |  connected   |
8     |          |   false   | vrf-created | not-provisioned | cfg-ready |
+------+--------+--------------+--------------------+--------------+--------------
+----------+----------+----------+-------------+-----------------+-----------+
Vrf Details
```

```
--- Time Elapsed: 93.298864ms ---
abc@abc-virtual-machine:~/GoDCApp/GoCommon/bin$


efa tenant epg create --tenant "t1" --name "epg1" --type extension --switchport-mode
trunk --single-homed-bfd-session-type auto --po po1 --vrf vrf1 --ctag-range 19 --l3-vni
5001 --anycast-ip 19:3.33.3.3/24 --bridge-domain 19:Auto-BD-2 --ctag-description
"19:Tenant L3 Extended BD" --l2-vni 19:2 --ip-mtu 19:1600 --ip-icmp-redirect 19:true --
ipv6-icmp-redirect 19:true
Error :     Input anycast ipv6 address configuration is needed for the ctag 19 when the
ipv6 icmp redirect configuration is input for the same ctag


efa tenant epg create --tenant "t1" --name "epg1" --type extension --switchport-mode
trunk --single-homed-bfd-session-type auto --po po1 --vrf vrf1 --ctag-range 19 --l3-vni
5001 --anycast-ip 19:3.33.3.3/24  --ctag-description "19:Tenant L3 Extended BD" --l2-vni
19:2 --ip-mtu 19:1600 --ip-icmp-redirect 19:true --ipv6-icmp-redirect 19:true --anycast-
ipv6 19:500::10/31
Error :     ICMP redirect feature is not supported on the device 10.20.246.1 with the
platform SLX9740-40C. It is supported on the SLX-9540 and SLX-9640 platforms only.


efa tenant epg create --tenant "t1" --name "epg1" --type extension --switchport-mode
trunk --single-homed-bfd-session-type auto --po po1 --vrf vrf1 --ctag-range 19 --l3-vni
5001 --anycast-ip 19:3.33.3.3/24  --ctag-description "19:Tenant L3 Extended BD" --l2-vni
19:2 --ip-mtu 19:1600 --ip-icmp-redirect 19:true --ipv6-icmp-redirect 19:true --anycast-
ipv6 19:500::10/31
    Device: 10.20.246.1
        Network Policy Error: VE configuration failed due to netconf rpc [error]
'%Error:IP address is not configured. ',
Error :     EndpointGroup Creation failed


efa tenant epg create --tenant "t1" --name "epg1" --type extension --switchport-mode
trunk --single-homed-bfd-session-type auto --po po1 --vrf vrf1 --ctag-range 19 --l3-vni
5001 --anycast-ip 19:3.33.3.3/24  --ctag-description "19:Tenant L3 Extended BD" --l2-vni
19:2 --ip-mtu 19:1600 --ip-icmp-redirect 19:true --ipv6-icmp-redirect 19:true --anycast-
ipv6 19:500::10/31 --local-ipv6 19,10.20.246.1:700::10/31 --local-ipv6
19,10.20.246.2:700::10/31

EndpointGroup created successfully.

--- Time Elapsed: 26.66300489s ---


efa tenant epg show --detail
================================================================================================
================================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19
           : VRF                          : vrf1
           : L3Vni                        : 5001

+------------+-------------+-----------+------------+-------------+
```

```
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
Port Property States

+------+-----------------------+-------+---------+--------------+--------------+----------
+-------------------------+--------------+--------+---------+----------------
+--------------+------------+------------+
| Ctag |         Ctag          | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 | Suppress
|         Local IP         | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND     |   IPv6
ND    |  Dev State  |  App State  |
|      |      Description      |       |         |              |              |  ARP/ND
|  [Device-IP->Local-IP]   |   IPv4/IPv6   |        |   MTU   | Managed Config | Other
Config |             |             |
+------+-----------------------+-------+---------+--------------+--------------+----------
+-------------------------+--------------+--------+---------+----------------
+--------------+------------+------------+
|  19  | Tenant L3 Extended BD |   2   |         | 3.33.3.3/24  |  500::10/31  |   T/T
| 10.20.246.1->700::10/31 |      T/T      |  1600  |         |     false      |
false     | provisioned | cfg-in-sync |
|      |                       |       |         |              |              |
| 10.20.246.2->700::10/31 |              |        |         |
|              |            |            |
+------+-----------------------+-------+---------+--------------+--------------+----------
+-------------------------+--------------+--------+---------+----------------
+--------------+------------+------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+-------------+---------------+-------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+-------------+---------------+-------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+------------------------------
+------------------------------------------------------------
+-----------------------------------------------+---------------------------------------
+
| Ctag |          AddressIP           |      AddressIPv6 :  Device-IP->[{Address-
IPv6,Vrf}] OR       |                GatewayIP :                  |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |       Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]        | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-
IP->{InfType,InfName,Gateway-IPv6} |
|      |
|                                                             |        Device-IP-
>{InfType,InfName}         |                                       |
+------+------------------------------
+------------------------------------------------------------
+-----------------------------------------------+---------------------------------------
+
```

```
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
================================================================

--- Time Elapsed: 146.093823ms ---
abc@abc-virtual-machine:~/GoDCApp/GoCommon/bin$
```

On SLX1:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

On SLX2:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

```
========================== EPG UPDATE - VRF DELETE=================================
efa tenant epg update --tenant "t1" --name "epg1" --operation vrf-delete --vrf vrf1

EndpointGroup updated successfully.

--- Time Elapsed: 11.522121773s ---
efa tenant epg show --detail
================================================================================================
================================================================
Name           : epg1
Tenant         : t1
Type           : extension
State          : epg-with-port-group-and-ctag-range
Description    :

Ports          :
POs            : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
Port Property States
```

```
+------+-------------------------+-------+---------+--------------+--------------
+----------+-----------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
Suppress |       Local IP        | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND     |
IPv6 ND    |  Dev State  |  App State  |
|      |       Description       |       |         |              |              |
ARP/ND  | [Device-IP->Local-IP] |   IPv4/IPv6   |        |   MTU   | Managed Config |
Other Config |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-----------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
|  19  | Tenant L2 Extended VLAN |   2   |         |              |              |
F/F    |                       |      T/T      |        |         |     false      |
false     | provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-----------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+--------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------
+
| Ctag |           AddressIP           |      AddressIPv6 :  Device-IP->[{Address-
IPv6,Vrf}] OR       |               GatewayIP :                   |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |       Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]        | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-
IP->{InfType,InfName,Gateway-IPv6} |
|      |
|                                                                 |        Device-IP-
>{InfType,InfName}         |                                                |
+------+--------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------
+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
=====================================================================
```

```
--- Time Elapsed: 125.742303ms ---
```

On SLX1:	On SLX2:
how runn int ve % No entries found.	show runn int ve % No entries found.

```
========================= EPG UPDATE - VRF ADD=================================

efa tenant epg update --tenant "t1" --name "epg1" --operation vrf-add --vrf vrf1 --ctag-
range 19 --l3-vni 5001 --anycast-ip 19:3.33.3.3/24  --ip-mtu 19:1600 --ip-icmp-redirect
19:true --ipv6-icmp-redirect 19:true --anycast-ipv6 19:500::10/31 --local-ipv6
19,10.20.246.1:700::10/31 --local-ipv6 19,10.20.246.2:700::10/31

EndpointGroup updated successfully.

--- Time Elapsed: 26.989502751s ---
efa tenant epg show --detail
================================================================================================
=====================================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19
          : VRF                           : vrf1
          : L3Vni                         : 5001

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
Port Property States

+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
Suppress |         Local IP         | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND
|   IPv6 ND    |  Dev State  |  App State  |
|      |       Description       |       |         |              |              |
ARP/ND  |  [Device-IP->Local-IP]   |   IPv4/IPv6   |        |   MTU   | Managed Config |
Other Config |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
|  19  | Tenant L3 Extended VLAN |   2   |         | 3.33.3.3/24  |  500::10/31  |
T/T    | 10.20.246.1->700::10/31 |      T/T      |  1600  |         |     false      |
false     | provisioned | cfg-in-sync |
|      |                         |       |         |              |
```

```
|          | 10.20.246.2->700::10/31 |               |        |         |
|              |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+--------------------------+--------------+--------+---------+----------------
+--------------+-------------+-------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+-------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+-------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+---------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------------
+
| Ctag |           AddressIP           |      AddressIPv6 :  Device-IP->[{Address-
IPv6,Vrf}] OR       |               GatewayIP :                  |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |       Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]        | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-
IP->{InfType,InfName,Gateway-IPv6} |
|      |
|                                                                 |        Device-IP-
>{InfType,InfName}         |                                                |
+------+---------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------------
+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
=============================================================================
```

```
--- Time Elapsed: 113.554904ms ---
```

On SLX1:	On SLX2:
show runn int ve interface Ve 19 vrf forwarding vrf1 ip anycast-address 3.33.3.3/24 ip mtu 1600 **ip icmp redirect** ipv6 anycast-address 500::10/31 ipv6 address 700::10/31 **ipv6 icmpv6 redirect** no shutdown ! interface Ve 8192 vrf forwarding vrf1 ipv6 address use-link-local-only no shutdown !	show runn int ve interface Ve 19 vrf forwarding vrf1 ip anycast-address 3.33.3.3/24 ip mtu 1600 **ip icmp redirect** ipv6 anycast-address 500::10/31 ipv6 address 700::10/31 **ipv6 icmpv6 redirect** no shutdown ! interface Ve 8192 vrf forwarding vrf1 ipv6 address use-link-local-only no shutdown !

```
========================== EPG UPDATE - CTAG ADD==================================
efa tenant epg update --tenant "t1" --name "epg1" --operation ctag-range-add --ctag-range
20 --anycast-ip 20:4.33.3.3/24  --ctag-description "20:Tenant L3 Extended BD" --l2-vni
20:3  --ip-icmp-redirect 20:true

EndpointGroup updated successfully.

--- Time Elapsed: 19.783074534s ---
efa tenant epg show --detail
================================================================================================
=============================================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19-20
          : VRF                          : vrf1
          : L3Vni                        : 5001

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
Port Property States

+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
```

```
Suppress |         Local IP         | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND
|   IPv6 ND    |  Dev State  |  App State  |
|      |       Description       |       |         |             |              |
ARP/ND  |  [Device-IP->Local-IP]   |   IPv4/IPv6   |        |   MTU   | Managed Config |
Other Config |             |             |
+------+-------------------------+-------+---------+-------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
|  19  | Tenant L3 Extended VLAN |   2   |         | 3.33.3.3/24  |  500::10/31  |
T/T    | 10.20.246.1->700::10/31 |      T/T      |  1600  |         |     false      |
false     | provisioned | cfg-in-sync |
|      |                         |       |         |              |
|          | 10.20.246.2->700::10/31 |               |        |         |
|              |             |             |
+------+-------------------------+-------+---------+-------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
|  20  |  Tenant L3 Extended BD  |   3   |         | 4.33.3.3/24  |              |
T/F    |                         |      T/F      |        |         |     false
|    false     | provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+-------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+---------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------
+
| Ctag |           AddressIP           |      AddressIPv6 :  Device-IP->[{Address-
IPv6,Vrf}] OR       |               GatewayIP :                  |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |       Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]        | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-
IP->{InfType,InfName,Gateway-IPv6} |
|      |
|                                                                 |        Device-IP-
>{InfType,InfName}         |                                                |
+------+---------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------
+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
=====================================================================
```

```
--- Time Elapsed: 133.454201ms ---
```

On SLX1:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 ip icmp redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

On SLX2:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 ip icmp redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

```
========================= EPG UPDATE - NETWORK PROPERTY
UPDATE==================================
efa tenant epg update --tenant "t1" --name "epg1" --operation network-property-update --
ip-icmp-redirect 20:false

EndpointGroup updated successfully.

--- Time Elapsed: 10.939994832s ---
efa tenant epg show --detail
================================================================================================
=====================================================================
Name           : epg1
Tenant         : t1
Type           : extension
State          : epg-with-port-group-and-ctag-range
Description    :

Ports          :
POs            : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19-20
          : VRF                          : vrf1
          : L3Vni                        : 5001

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
```

```
Port Property States

+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+---------+---------+----------------
+--------------+-------------+--------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
Suppress |         Local IP         | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND
|   IPv6 ND    |  Dev State  |  App State  |
|      |       Description       |       |         |              |              |
ARP/ND  |  [Device-IP->Local-IP]   |   IPv4/IPv6   |        |   MTU   | Managed Config |
Other Config |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+---------+---------+----------------
+--------------+-------------+--------------+
|  19  | Tenant L3 Extended VLAN |   2   |         | 3.33.3.3/24  |  500::10/31  |
T/T    | 10.20.246.1->700::10/31 |      T/T      |  1600  |         |     false      |
false     | provisioned | cfg-in-sync |
|      |                         |       |         |              |
|          | 10.20.246.2->700::10/31 |               |        |         |
|              |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+---------+---------+----------------
+--------------+-------------+--------------+
|  20  |  Tenant L3 Extended BD  |   3   |         | 4.33.3.3/24  |              |
T/F    |                         |      F/F      |        |         |     false
|    false     | provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+---------+---------+----------------
+--------------+-------------+--------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+---------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------
+
| Ctag |           AddressIP           |      AddressIPv6 :  Device-IP->[{Address-
IPv6,Vrf}] OR       |               GatewayIP :                   |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |       Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]        | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-
IP->{InfType,InfName,Gateway-IPv6} |
|      |
|                                                                 |        Device-IP-
>{InfType,InfName}         |                                                |
+------+---------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------
+
DHCP Relay Ips
```

```
For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
=====================================================================

--- Time Elapsed: 194.609986ms ---
```

On SLX1:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

On SLX2:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 20
vrf forwarding vrf1
ip anycast-address 4.33.3.3/24
no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

```
========================== EPG UPDATE - NETWORK PROPERTY
ADD==================================
efa tenant epg update --tenant "t1" --name "epg1" --operation network-property-add --ip-
icmp-redirect 20:true

EndpointGroup updated successfully.

--- Time Elapsed: 8.91347074s ---
efa tenant epg show --detail
================================================================================================
=====================================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19-20
           : VRF                          : vrf1
           : L3Vni                        : 5001

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
```

```
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
Port Property States

+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
Suppress |         Local IP         | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND
|   IPv6 ND    |  Dev State  |  App State  |
|      |       Description       |       |         |              |              |
ARP/ND  |  [Device-IP->Local-IP]   |   IPv4/IPv6   |        |   MTU   | Managed Config |
Other Config |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
|  19  | Tenant L3 Extended VLAN |   2   |         | 3.33.3.3/24  |  500::10/31  |
T/T    |  10.20.246.1->700::10/31 |      T/T      |  1600  |         |     false
|    false     | provisioned | cfg-in-sync |
|      |                         |       |         |              |
|          |  10.20.246.2->700::10/31 |               |        |         |
|              |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
|  20  |  Tenant L3 Extended BD  |   3   |         | 4.33.3.3/24  |              |
T/F    |                         |      T/F      |        |         |     false
|    false     | provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+----------+-------------------------+---------------+--------+---------+----------------
+--------------+-------------+--------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+--------------------------------
+-------------------------------------------------
+---------------------------------------------+-----------------------------------------
+
| Ctag |           AddressIP           | AddressIPv6 :  Device-IP->[{Address-IPv6,Vrf}]
OR  |               GatewayIP :                  |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |  Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]   | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-IP-
>{InfType,InfName,Gateway-IPv6} |
|      |
|                                                    |        Device-IP-
>{InfType,InfName}         |                                          |
+------+--------------------------------
+-------------------------------------------------
+---------------------------------------------+-----------------------------------------
+
```

```
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================
================================================================

--- Time Elapsed: 129.081769ms ---
abc@abc-virtual-machine:~/GoDCApp/GoCommon/bin$
```

On SLX1:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 ip icmp redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

On SLX2:

```
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 ipv6 icmpv6 redirect
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 ip icmp redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
```

```
========================= EPG UPDATE - NETWORK PROPERTY
DELETE==================================

efa tenant epg update --tenant "t1" --name "epg1" --operation network-property-delete --
ipv6-icmp-redirect 19:true

EndpointGroup updated successfully.

--- Time Elapsed: 8.086847111s ---
efa tenant epg show --detail
================================================================================================
================================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po1
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 19-20
          : VRF                         : vrf1
          : L3Vni                       : 5001

+------------+------------+-----------+------------+-------------+
```

```
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po1  | provisioned | cfg-in-sync |
+------+-------------+-------------+
Port Property States

+------+-------------------------+-------+---------+--------------+--------------
+----------+--------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
Suppress |         Local IP         | Icmp Redirect | IP MTU | IPv6 ND |    IPv6 ND
|   IPv6 ND    |  Dev State  |  App State  |
|      |       Description       |       |         |              |              |
ARP/ND  |  [Device-IP->Local-IP]   |   IPv4/IPv6   |        |   MTU   | Managed Config |
Other Config |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+--------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
|  19  | Tenant L3 Extended VLAN |   2   |         | 3.33.3.3/24  |  500::10/31  |
T/T    | 10.20.246.1->700::10/31  |      T/F      |  1600  |         |     false
|    false     | provisioned | cfg-in-sync |
|      |                         |       |         |              |              |
|          | 10.20.246.2->700::10/31  |               |        |         |
|              |             |             |
+------+-------------------------+-------+---------+--------------+--------------
+----------+--------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
|  20  |  Tenant L3 Extended BD  |   3   |         | 4.33.3.3/24  |              |
T/F    |                          |      T/F      |        |         |     false
|    false     | provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+----------+--------------------------+---------------+--------+---------+----------------
+--------------+-------------+-------------+
Network Property [Flags    : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------
+-------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime | Config
Type |
+------+----------------+--------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

+------+------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+------------+-------------+-----------+------------+-------------+
Network Property ACLs

+------+-------------------------------
+------------------------------------------------------
+------------------------------------------+------------------------------------------
+
| Ctag |          AddressIP           |    AddressIPv6 :  Device-IP->[{Address-
IPv6,Vrf}] OR   |               GatewayIP :                 |
GatewayIPv6 :                |
|      | Device-IP->[{Address-IP,Vrf}] |     Device-IP->[{Address-
IPv6,Vrf,InfType,InfName}]    | Device-IP->{Gateway-IP,InfType,InfName} OR  | Device-IP-
>{InfType,InfName,Gateway-IPv6} |
|      |
```

```
|                                                        |        Device-IP-
>{InfType,InfName}         |                                             |
+------+------------------------------
+--------------------------------------------------------
+----------------------------------------------+----------------------------------------
+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

==========================================================================================
=====================================================================

--- Time Elapsed: 147.804073ms ---
```

<table>
<tr>
<td>On SLX1:
<pre>
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 ip icmp redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
</pre></td>
<td>On SLX2:
<pre>
show runn int ve
interface Ve 19
 vrf forwarding vrf1
 ip anycast-address 3.33.3.3/24
 ip mtu 1600
 ip icmp redirect
 ipv6 anycast-address 500::10/31
 ipv6 address 700::10/31
 no shutdown
!
interface Ve 20
 vrf forwarding vrf1
 ip anycast-address 4.33.3.3/24
 ip icmp redirect
 no shutdown
!
interface Ve 8192
 vrf forwarding vrf1
 ipv6 address use-link-local-only
 no shutdown
!
</pre></td>
</tr>
</table>

Update Anycast IP on an Existing Tenant Network

You can ad or delete an anycast IP (both IPv4 and IPv6) on an existing tenant network.

About This Task

Follow this procedure to add or delete anycast IPv4 and anycast IPv6 to or from an existing L3 EPG tenant network.

- You can provide IPv6 ND attributes (ipv6 nd mtu, ipv6 nd prefix, ipv6 nd m or o flags) along with anycast-ipv6 when you add an anycast IP.
- You can provide only one anycast-ipv4 and one anycast-ipv6 per tenant network even though SLX supports multiple anycast IPv4 or IPv6 per VE.

Typical usage of the API from Openstack Integration is as follows:

1. EPG create with a ctag.
2. EPG update port-group-add with endpoints (po or phy).
3. EPG update vrf-add with anycast-ipv4.
4. EPG update anycast-ip-add with anycast-ipv6.

Procedure

1. To add an anycast IP when you update a tenant EPG network, run the following command:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
        --operation <anycast-ip-add | anycast-ip-delete>
        --anycast-ip <ctag:anycast-ipv4> --anycast-ipv6 <ctag:anycast-ipv6>
```

   The following example configures an anycast IP:

```
efa tenant epg show --detail
================================================================================================
================================================================================
Name          : ten1epg1
Tenant        : vlanTen1
Description   :
Type          : extension
Ports         : 10.20.246.15[0/1]
              : 10.20.246.16[0/1]
POs           :
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11
              : vrf                 : ten1vrf1
              : l3-vni              : 8188
Network Property [Flags : * - Native Vlan]
+----+------+------------+-------+------+----------------------+--------------+--------
+---------------+---------------+-------------+-------------+
|Ctag|L2-Vni|Anycast-IPv4|Anycast| BD   | Local IP             |  Ctag-       | Mtu-
| Managed       | Other         |  Dev-state  |  App-state  |
|    |      |            |-IPv6  | -name| (Device-IP->Local-IP)|  Description | IPv6-
ND| Config-IPv6-ND| Config-IPv6-ND|             |             |
+----+------+------------+-------+------+----------------------+--------------+--------
+---------------+---------------+-------------+-------------+
| 11 |11    |10.0.11.1/24|       |      |                      | Tenant L3    |
| False         | False         | provisioned | cfg-in-sync |
|    |      |            |       |      |                      | Extended VLAN|
|               |               |             |                      |
+----+------+------------+-------+------+----------------------+--------------+--------
+---------------+---------------+-------------+-------------+
================================================================================================
============================================================

efa tenant epg update --name ten1epg1 --tenant vlanTen1 --operation anycast-ip-add --
anycast-ipv6 11:10::1/123

efa tenant epg show --detail
================================================================================================
============================================================
Name          : ten1epg1
Tenant        : vlanTen1
Description   :
Type          : extension
Ports         : 10.20.246.15[0/1]
              : 10.20.246.16[0/1]
POs           :
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11
              : vrf                 : ten1vrf1
              : l3-vni              : 8188
Network Property [Flags : * - Native Vlan]
+------+--------+--------------+--------------+---------
```

```
+--------------------------------+-------------------------+-------------
+-----------------------+---------------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-IPv4 | Anycast-IPv6 | BD-name | Local IP (Device-IP->Local-
IP) |    Ctag-Description     | Mtu-IPv6-ND | ManagedConfig-IPv6-ND | OtherConfig-IPv6-
ND |  Dev-state  |  App-state  |
+------+--------+--------------+--------------+---------
+--------------------------------+-------------------------+-------------
+-----------------------+---------------------+-------------+-------------+
| 11   | 11     | 10.0.11.1/24 | 10::1/123    |
|                                | Tenant L3 Extended VLAN |             |
False                  | False               | provisioned | cfg-in-sync |
+------+--------+--------------+------+---------+--------------------------------
+-------------------------+-------------+----------------------+--------------------
+-------------+-------------+
================================================================================================
=========================================================================
```

2. To delete an anycast IP on an existing EPG tenant network, run the following command:

```
efa tenant epg update --name ten1epg1 --tenant vlanTen1 --operation anycast-ip-delete
--anycast-ipv6 11:10::1/123

efa tenant epg show
================================================================================================
=========================================================================
Name           : ten1epg1
Tenant         : vlanTen1
Description    :
Type           : extension
Ports          : 10.20.246.15[0/1]
               : 10.20.246.16[0/1]
POs            :
Port Property  : switchport mode     : trunk
               : native-vlan-tagging : false
NW Policy      : ctag-range          : 11
               : vrf                 : ten1vrf1
               : l3-vni              : 8188
Network Property [Flags : * - Native Vlan]
+------+--------+--------------+--------------+---------
+--------------------------------+-------------+-----------------------
+----------------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-IPv4 | Anycast-IPv6 | BD-name | Local IP (Device-IP->Local-
IP) | Mtu-IPv6-ND | ManagedConfig-IPv6-ND | OtherConfig-IPv6-ND |  Dev-state  |  App-
state  |
+------+--------+--------------+--------------+---------
+--------------------------------+-------------+-----------------------
+----------------------+-------------+-------------+
| 11   | 11     | 10.0.11.1/24 |              |
|                                |             | False                 |
False                | provisioned | cfg-in-sync |
+------+--------+--------------+--------------+---------
+--------------------------------+-------------+-----------------------
+----------------------+-------------+-------------+
For 'unstable' entities, run 'efa tenant po/vrf show' for details
================================================================================================
=========================================================================
```

Configure Multiple Anycast IP

You can provide multiple anycast IP address for each tenant ctag when you create or update an EPG, add ctag-range, add VRF, and add or delete anycast IP. The multiple anycast IP address is configured under the `interface Ve` of the switching hardware.

About This Task

Follow this procedure to configure multiple anycast IP.

Procedure

1. To configure multiple anycast IP when you create an EPG, run the following command:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>

        --switchport-mode <mode>
        --port <port-list> --po <po-list>

        --ctag-range <ctag-range>
        --vrf <vrf-name>

        --anycast-ip <value> --anycast-ip <value>
        --anycast-ipv6 <value> --anycast-ipv6 <value>
```

2. To configure multiple anycast IP when you update an EPG, run the following command:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>

        --operation vrf-add | ctag-range-add |anycast-ip-add | anycast-ip-delete|
        --ctag-range <ctag-range> --vrf  <vrf-name>

        --anycast-ip <value> --anycast-ip <value>
        --anycast-ipv6 <value> --anycast-ipv6 <value>
```

Example

The following example configures multiple anycast IP when you create an EPG:

```
efa tenant epg create --name e1 --tenant tenant11 --po po1 --switchport-mode trunk-no-
default-native --vrf v1 --ctag-range 101-102  --anycast-ip 101:1.1.1.254/24 --anycast-ip
101:2.1.1.254/24 --anycast-ip 101:3.1.1.254/24 --anycast-ipv6 101:1::1/124 --anycast-ipv6
102:2::1/124
efa tenant epg show -detail
================================================================================
Name          : e1
Tenant        : tenant11
Type          : extension
State         : epgf-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po1
Port Property : SwitchPort Mode               : trunk-no-default-native
              : Single-Homed BFD Session Type : Auto
NW Policy     : Ctag Range                    : 101-103
              : VRF                           : v1
              : L3Vni                         : 8192

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUP | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
```

```
POrt Property ACLs

+------+--------------+-------------+
| Port |  Dev State   |  App State  |
+------+--------------+-------------+
| po1  | provisionied | cfg-in-sync |
+------+--------------+-------------+
Port Property States

+-----+-------------+-----+----+-------------+--------+--------+-------------------------
+---+------+--------------+------------+-----------+-----------+
| Ctag|   Ctag      |L2Vni|BD  |Anycast IPv4 |Anycast |Suppress|        Local IP
|IP |IPv6  |IPv6 ND       |IPv6 ND     | Dev State | App State |
|     | Description |     |Name|             |IPv6    | ARP/ND |  [Device-IP->Local-
IP] |MTU|MTU ND|Managed Oonfig|Other Config|           |           |
+-----+-------------+-----+----+-------------+--------+--------+-------------------------
+---+------+--------------+------------+-----------+-----------|
| 101 |  Tenant L3  |10000|    |2.1.1.254/24 |1::1/124|   T/T  |10.20.246.3-
>1.10.1.1/24|   |      |    false     |   false    |porvisioned|cfg-in-sync|
|     |Extended VLAN|     |    |3.1.1.254/24 |        |        |
|   |      |             |           |           |          |
|     |             |     |    |           |        |        |
|   |      |             |           |           |          |
|     |             |     |    |1.1.1.254/24 |        |        |
|   |      |             |           |           |          |
+-----+-------------+-----+----+-------------+--------+--------+-------------------------
+---+------+--------------+------------+-----------+-----------|
| 102 |  Tenant L3  |10001|    |             |2::1/124|   T/T  |
|   |      |    false     |   false    |porvisioned|cfg-in-sync|
|     |Extended VLAN|     |    |             |        |        |
|   |      |             |           |           |          |
+-----+-------------+-----+----+-------------+--------+--------+-------------------------
+---+------+--------------+------------+-----------+-----------+
Network Property [Flags : * - Native Vlan]

+----+--------------+------------+-------------+------------------+-----------+
|Ctag|IPv6 ND Prefix|No Advertise|Valid Lifetime|Preferred Lifetime|Config Type|
+----+--------------+------------+-------------+------------------+-----------+
IPv6 ND Prefix Flags

+------+----------------+-------------+-----------+------------+-------------+
| Ctag | MAC ACL Prefix | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------+----------------+-------------+-----------+------------+-------------+
Network Property ACLs
```

```
For 'unstable' entities, run 'efa tenant po/vrf show' for details
```

10.20.246.25

```
SLX# show running-config interface
Ve
interface Ve 101
 vrf forwarding v1
 ip anycast-address 1.1.1.254/24
 ip anycast-address 2.1.1.254/24
 ip anycast-address 3.1.1.254/24
 ipv6 anycast-address 1::1/124
 no shutdown
!
interface Ve 102
 vrf forwarding v1
 ipv6 anycast-address 2::1/124
 no shutdown
!
interface Ve 4090
 ip address 10.20.20.2/31
 no shutdown
!
interface Ve 5120
 vrf forwarding v1
 ipv6 address use-link-local-only
 no shutdown
!
```

10.20.246.26

```
SLX# show running-config interface
Ve
interface Ve 101
 vrf forwarding v1
 ip anycast-address 1.1.1.254/24
 ip anycast-address 2.1.1.254/24
 ip anycast-address 3.1.1.254/24
 ipv6 anycast-address 1::1/124
 no shutdown
!
interface Ve 102
 vrf forwarding v1
 ipv6 anycast-address 2::1/124
 no shutdown
!
interface Ve 4090
 ip address 10.20.20.3/31
 no shutdown
!
interface Ve 5120
 vrf forwarding v1
 ipv6 address use-link-local-only
 no shutdown
!
```

The following example configures multiple anycast IP when you update an EPG:

```
efa  tenant epg update --name e1 --tenant tenant11 --operation anycast-ip-add --anycast-
ip 101:4.1.1.1/24
```

10.20.246.25

```
SLX# show running-config interface
Ve
interface Ve 101
 vrf forwarding v1
 ip anycast-address 1.1.1.254/24
 ip anycast-address 2.1.1.254/24
 ip anycast-address 3.1.1.254/24
 ip anycast-address 4.1.1.1/24
 ipv6 anycast-address 1::1/124
 no shutdown
!
interface Ve 102
 vrf forwarding v1
 ipv6 anycast-address 2::1/124
 no shutdown
!
interface Ve 4090
 ip address 10.20.20.2/31
 no shutdown
!
interface Ve 5120
 vrf forwarding v1
 ipv6 address use-link-local-only
 no shutdown
!
```

10.20.246.26

```
SLX# show running-config interface
Ve
interface Ve 101
 vrf forwarding v1
 ip anycast-address 1.1.1.254/24
 ip anycast-address 2.1.1.254/24
 ip anycast-address 3.1.1.254/24
 ip anycast-address 4.1.1.1/24
 ipv6 anycast-address 1::1/124
 no shutdown
!
interface Ve 102
 vrf forwarding v1
 ipv6 anycast-address 2::1/124
 no shutdown
!
interface Ve 4090
 ip address 10.20.20.3/31
 no shutdown
!
interface Ve 5120
 vrf forwarding v1
 ipv6 address use-link-local-only
 no shutdown
!
```

Configure IPv6 Neighbor Discovery (ND) on a Tenant Network

You can configure IPv6 neighbor discovery (ND) attributes (MTU, M flag, O flag, and Prefixes) for each tenant network (ctag).

About This Task

Follow this procedure to configure IPv6 ND attributes when you create or update an EPG or add ctag-range, VRF, and anycast IP.

Procedure

1. To configure IPv6 ND when you create an EPG, run the following command:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>
     --ipv6-nd-mtu <ipv6-mtu>
     --ipv6-nd-managed-config <true | false>
     --ipv6-nd-other-config <true | false>
     --ipv6-nd-prefix <ctag:list-of-prefix>
     --ipv6-nd-prefix-valid-lifetime <ctag,prefix:validTime>
     --ipv6-nd-prefix-preferred-lifetime <ctag,prefix:preferredTime>
     --ipv6-nd-prefix-no-advertise <ctag,prefix:noadvertiseflag>
     --ipv6-nd-prefix-config-type <ctag,prefix:configType(no-autoconfig| no-onlink |
off-link)>
```

2. To configure IPv6 ND when you update an EPG, run the following command:

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
      --operation <ctag-range-add | vrf-add | anycast-ip-add> --ctag-range <ctag-
range> --vrf <vrf-name>
      --anycast-ip <ctag:anycast-ip> --anycast-ipv6 <ctag:anycast-ipv6>
      --ipv6-nd-mtu <ipv6-mtu>
      --ipv6-nd-managed-config <true | false>
      --ipv6-nd-other-config <true | false>
      --ipv6-nd-prefix <ctag:list-of-prefix>
      --ipv6-nd-prefix-valid-lifetime <ctag,prefix:validTime>
      --ipv6-nd-prefix-preferred-lifetime <ctag,prefix:preferredTime>
      --ipv6-nd-prefix-no-advertise <ctag,prefix:noadvertiseflag>
      --ipv6-nd-prefix-config-type <ctag,prefix:configType(no-autoconfig| no-onlink |
off-link)>
```

The following example configures IPv6 Neighbor Discovery (ND) on a Tenant Network:

```
efa tenant show
+----------+---------+------------+-------------+-------------+-----------+-----------
+----------------------+
|   Name   |  Type   | VLAN Range | L2VNI Range | L3VNI Range | VRF Count | Enable
BD |        Ports         |
+----------+---------+------------+-------------+-------------+-----------+-----------
+----------------------+
| vlanTen1 | private |   11-20    |             |             |    10     |   false
| 10.20.246.15[0/1-10] |
|          |         |            |             |             |           |
| 10.20.246.16[0/1-10] |
+----------+---------+------------+-------------+-------------+-----------+-----------
+----------------------+


efa tenant vrf show
+---------+---------+-----------+-----------+-------------+-----+------+-------
+------------------+------------+------------+
|  Name   | Tenant  |Routing    |Centralized| Redistribute| Max
| Local| Enable|      State       |  Dev State |  App State |
|         |         |Type       |Routers    |             | Path|
```

```
Asn  | GR    |                  |            |            |
+---------+---------+-----------+-----------+-------------+-----+------+-------
+------------------+------------+------------+
|ten1vrf1 |vlanTen1 |distributed|           |  connected  | 8
|      | false |vrf-device-created| provisioned| cfg-in-sync|
+---------+---------+-----------+-----------+-------------+-----+------+-------
+------------------+------------+------------+


efa tenant epg create --name ten1epg1 --tenant vlanTen1
--port 10.20.246.15[0/1],10.20.246.16[0/1] --switchport-mode trunk --ctag-range 11-13
    --anycast-ip 11:10.0.11.1/24 --anycast-ip 12:10.0.12.1/24 --anycast-
ip 13:10.0.13.1/24
    --ipv6-nd-mtu 12:1600 --ipv6-nd-managed-config 12:true --ipv6-nd-
other-config 12:true
    --ipv6-nd-prefix 12:1:5::/64 --ipv6-nd-prefix-valid-lifetime 12,1:5::/64:2000
--ipv6-nd-prefix-preferred-lifetime 12,1:5::/64:2000
    --ipv6-nd-prefix 12:1:6::/64 --ipv6-nd-prefix-valid-lifetime 12,1:6::/64:2001
--ipv6-nd-prefix-preferred-lifetime 12,1:6::/64:2001
    --vrf ten1vrf1

efa tenant epg show
+---------+---------+----------+-------------------+---+------------+------------
+------+---------+------+------+
|  Name   | Tenant  |   Type   |       Ports       |PO
| SwitchPort | Native Vlan| Ctag |   Vrf   | L3Vni| State|
|         |         |          |                   |   |    Mode
|   Tagging  | Range|         |      |      |
+---------+---------+----------+-------------------+---+------------+------------
+------+---------+------+------+
|ten1epg1 |vlanTen1 | extension| 10.20.246.15[0/1] |   |   trunk
|    false   | 11-13| ten1vrf1| 8192 |      |
|         |         |          | 10.20.246.16[0/1] |   |
|            |      |         |      |      |
+---------+---------+----------+-------------------+---+------------+------------
+------+---------+------+------+


efa tenant epg show --detail

================================================================================
Name          : ten1epg1
Tenant        : vlanTen1
Type          : extension
State         :
Description   :
Ports         : 10.20.246.15[0/1]
              : 10.20.246.16[0/1]
POs           :
Port Property : SwitchPort Mode     : trunk
              : Native Vlan Tagging : false
NW Policy     : Ctag Range          : 11-13
              : VRF                 : ten1vrf1
              : L3Vni               : 8192
+----+-------------+-----+----+------------+-------+----------------------+------
+--------------+------------+-----------+-----------+
|Ctag|    Ctag     |L2Vni|BD  |Anycast IPv4|Anycast|      Local
IP        |IPv6  |IPv6 ND       |  IPv6 ND   | Dev State | App State |
|    |Description  |     |Name|            |IPv6   |[Device-IP->Local-
IP] |ND Mtu|Managed Config|Other Config|           |           |
+----+-------------+-----+----+------------+-------+----------------------+------
+--------------+------------+-----------+-----------+
| 11 |Tenant L3    | 11  |    |10.0.11.1/24|
|                     |     |     false    |   false    |provisioned|cfg-in-sync|
```

```
|    |Extended VLAN|     |    |            |       |
|      |              |           |          |           |           |
+----+-------------+-----+----+------------+-------+----------------------+------
+--------------+------------+-----------+-----------+
| 12 | Tenant L3   | 12  |    |10.0.12.1/24|
|                    |1600  |     true     |   true     |provisioned|cfg-in-sync|
|    |Extended VLAN|     |    |            |       |
|      |              |           |          |           |           |
+----+-------------+-----+----+------------+-------+----------------------+------
+--------------+------------+-----------+-----------+
| 13 | Tenant L3   | 13  |    |10.0.13.1/24|
|                     |      |     false    |   false    |provisioned|cfg-in-sync|
|    |Extended VLAN|     |    |            |       |
|      |              |           |          |           |           |
+----+-------------+-----+----+------------+-------+----------------------+------
+--------------+------------+-----------+-----------+
Network Property [Flags : * - Native Vlan]
+-----+--------+----------+---------+----------+-------+
| Ctag| IPv6 ND| No       | Valid   | Preferred| Config|
|     | Prefix | Advertise| Lifetime| Lifetime | Type  |
+-----+--------+----------+---------+----------+-------+
|  12 |1:5::/64| false    |  2000   |  2000    |       |
+-----+--------+----------+---------+----------+-------+
|  12 |1:6::/64| false    |  2001   |  2001    |       |
+-----+--------+----------+---------+----------+-------+
IPv6 ND Prefix Flags
For 'unstable' entities, run 'efa tenant po/vrf show' for details
==================================================================
```

Configure BFD Session Type for an Endpoint Group

You can determine the session type for a Bidirectional Forwarding Detection (BFD) session formed over a Cluster Edge Port (CEP) port.

About This Task

You can assign a session type as you create an endpoint group. You can also assign a session type to an existing endpoint group.

The default is auto, which means that the BFD session type is automatically determined based on whether the service type is set to extension (software) or Layer 3 hand-off (hardware).

The value of `--single-homed-bfd-session-type` is configured for one endpoint group and then propagated to all Ethernet and single-homed port channel interfaces defined for that endpoint group.

XCO does not distinguish between SRIOV (single-root input/output virtualization) and non-SRIOV connections. Therefore, it treats both connections the same way. If you want to use hardware-based BFD sessions for CEP non-SRIOV connections, then create an endpoint group that contains all the CEP non-SRIOV connections and set the `--single-homed-bfd-session-type` to `hardware`.

Procedure

1. To assign a BFD session type when you create an endpoint group, run the following command:

```
$ efa tenant epg create --name epg5 --tenant tenant11 --port 10.20.216.15[0/11]
,10.20.216.16[0/11] --po po1 --switchport-mode trunk --single-homed-bfd-session-type
auto
```

   In this example, the session type is set to 'auto'.
2. To assign a BFD session type to an existing endpoint group, run the following command:

```
$ efa tenant epg update --name epg5 --tenant tenant11 --operation port-group-add
--port 10.20.216.15[0/11],10.20.216.16[0/11] --po po1 --switchport-mode trunk
--single-homed-bfd-session-type hardware
```

   In this example, the session type is set to 'hardware'.

Configure Cluster Edge Port (CEP) Cluster Tracking for Endpoint Groups

XCO does not provision `reload-delay 90` configuration on Cluster Edge Port (CEP) interfaces. XCO instead provisions `cluster-track` configuration on CEP interfaces when the CEP is configured as a member of an EPG (endpoint group) during the creation or update of endpoint groups.

About This Task

Follow this procedure to configure Cluster Edge Port (CEP) cluster tracking for an endpoint group (EPG).

Procedure

To configure CEP Cluster Tracking for an EPG, run the following command:

```
# efa tenant po create --name po11 --tenant tenant1 -speed 10Gbps --negotiation active
--port 10.24.80.134[0/15] => CEP # efa tenant po create --name po12 --tenant tenant1
--speed 10Gbps --negotiation active --port 10.24.80.134[0/25],10.24.80.135[0/25] => CCEP

# efa tenant epg create --name ten1epg1 --tenant tenant1 --switchport-mode trunk --ctag-
range 1001
--port 10.24.80.134[0/35] --po po11,po12
```

Enable Cluster Tracking on CEP Interfaces

By default, XCO enables reload-delay configuration on all Cluster Edge Port (CEP) Interfaces. Reload-delay and cluster-tracking configurations are mutually exclusive.

About This Task

When cluster tracking is enabled, an interface can track the state of an MCT (multi-chassis tunnel) cluster and divert traffic to alternative paths when a cluster is down for reasons such as maintenance mode.

To enable cluster tracking on CEP interfaces, you must remove the reload-delay configuration.

Procedure

1. Run the following command to remove the reload-delay configuration on CEP interface:

```
# efa inventory device execute-cli --ip 10.18.120.187 --command "Interface ethernet
0/1, no reload-delay enable" --config
```

2. Run the following command to configure cluster tracking:

```
# efa inventory device execute-cli --ip 10.18.120.187 --command "Interface ethernet
0/1, cluster-track" --config
```

Configure Suppress Address Resolution Protocol and Neighbor Discovery on VLAN or Bridge Domain

You can configure suppress address resolution protocol and neighbor discovery on VLAN or Bridge Domain.

Before You Begin

Provide an option to enable or disable suppress ARP (suppress address resolution protocol) or ND (neighbor discovery) at the tenant network level so that you can choose to enable suppress ARP or ND per tenant network.

This option is mainly useful for a single rack small data center deployment where the suppress ARP or ND configuration on tenant network is not needed.

About This Task

Follow this procedure to configure suppress address resolution protocol and neighbor discovery on VLAN or Bridge Domain

Note

- XCO configures the suppress ARP or ND on the VLAN or BD (Bridge Domain) associated with the L3 tenant networks belonging to the distributed router.
- XCO does not configure the suppress ARP or ND on the VLAN or BD associated with the L3 tenant networks belonging to the centralized router.
- XCO doesn't configure the suppress ARP or ND on the VLAN or BD associated with the L2 tenant networks.
- For all the L3 tenant networks belonging to the distributed router, suppress ARP or ND is displayed as "true" after the upgrade from 2.4.x to 2.5.5.
- For all the L3 tenant networks belonging to the centralized router, suppress ARP or ND is displayed as "false" after the upgrade from 2.4.x to 2.5.5.
- For all the L2 tenant networks, suppress ARP or ND is displayed as "false" after the upgrade from 2.4.x to 2.5.5.

Procedure

To configure suppress ARP or ND for performance tuning when you create a VLAN or BD, run the following command:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>
       --port <port-list> --switchport-mode <trunk|access> --ctag-range <ctag-range>
```

```
       --anycast-ip <ctag:anycast-ipv4> --anycast-ipv6 <ctag:anycast-ipv6> --vrf ten1vrf1
       --suppress-arp <ctag:true|false> --suppress-nd <ctag:true|false>
 efa tenant epg update --name <epg-name> --tenant <tenant-name> --operation vrf-add
        --anycast-ip <ctag:anycast-ipv4> --anycast-ipv6 <ctag:anycast-ipv6>  --vrf
ten1vrf1
        --suppress-arp <ctag:true|false> --suppress-nd <ctag:true|false>
efa tenant epg update --name <epg-name> --tenant <tenant-name> --operation ctag-range-add
       --ctag-range <ctag-range> --anycast-ip <ctag:anycast-ipv4> --anycast-ipv6
<ctag:anycast-ipv6>
       --suppress-arp <ctag:true|false> --suppress-nd <ctag:true|false>
 efa tenant epg update --name <epg-name> --tenant <tenant-name> --operation anycast-ip-
add
      --anycast-ip <ctag:anycast-ipv4> --anycast-ipv6 <ctag:anycast-ipv6>
        --suppress-arp <ctag:true|false> --suppress-nd <ctag:true|false>
efa fabric show
Fabric Name: default, Fabric Description: Default Fabric, Fabric Stage: 3, Fabric Type:
clos, Fabric Status: created
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE | DEVICE STATE | APP STATE | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+

Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric Status:
settings-updated

Updated Fabric Settings: BGP-LL

+-------------+---+-----------+-------+------+---------------------+-------------
+-------------------+-----------------------------------------------+---------+-------+
| IP ADDRESS  |POD| HOST NAME | ASN   | ROLE | DEVICE STATE        | APP STATE   | CONFIG
GEN REASON | PENDING CONFIGS                          | VTLB ID | LB ID |
+-------------+---+-----------+-------+------+---------------------+-------------
+-------------------+-----------------------------------------------+---------+-------+
| 10.20.246.1 |   | SLX-1     | 64512 | Spine| provisioned         | cfg in-sync |
NA                | NA                                          | NA      | 1     |
| 10.20.246.7 |   | SLX       | 65000 | Leaf | provisioning failed | cfg ready   |
IA,IU,MD,DA       | SYSP-C,MCT-C,MCT-PA,BGP-C,INTIP-C,EVPN-C,O-C | 2       | 1     |
| 10.20.246.8 |   | slx-8     | 65000 | Leaf | provisioned         | cfg in-sync |
NA                | NA                                          | 2       | 1     |
+-------------+---+-----------+-------+------+---------------------+-------------
+-------------------+-----------------------------------------------+---------+-------+
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password , BFD-RX - Bfd Rx
Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd multiplier,
BFD-ENABLE - Enable Bfd, BGP-MULTIHOP - BGP ebgp multihop, P2PLR - Point-to-Point Link
Range, MCTLR - MCT Link Range, LOIP - Loopback IP Range


CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
```

```
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

efa tenant show
+------+---------+------------+-------------+-------------+-----------+-----------
+------------------------+
| Name |  Type   | VLAN Range | L2VNI Range | L3VNI Range | VRF Count | Enable BD
|         Ports          |
+------+---------+------------+-------------+-------------+-----------+-----------
+------------------------+
| ten1 | private |   11-20    |             |             |    10     |   false   |
10.20.246.15[0/1-10]  |
|      |         |            |             |             |           |           |
10.20.246.16[0/1-10]  |
+------+---------+------------+-------------+-------------+-----------+-----------
+------------------------+
| ten2 | private |   21-30    |             |             |    10     |   true    |
10.20.246.15[0/11-20] |
|      |         |            |             |             |           |           |
10.20.246.16[0/11-20] |
+------+---------+------------+-------------+-------------+-----------+-----------
+------------------------+
efa tenant vrf show
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+-------------+
|   Name   | Tenant | Routing Type | Centralized Routers | Redistribute | Max Path |
Local Asn | Enable GR |       State        |  Dev State  |  App State  |
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+-------------+
| ten1vrf1 |  ten1  | distributed  |                     |  connected   |    8
|           |   false   | vrf-device-created | provisioned | cfg-in-sync |
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+-------------+
| ten2vrf1 |  ten2  | distributed  |                     |  connected   |    8
|           |   false   | vrf-device-created | provisioned | cfg-in-sync |
+----------+--------+--------------+---------------------+--------------+----------
+-----------+-----------+--------------------+-------------+-------------+
efa tenant epg create --name ten1epg1 --tenant ten1 --port
10.20.246.15[0/1],10.20.246.16[0/1]
     --switchport-mode trunk --ctag-range 11-12
     --anycast-ip 11:10.0.11.1/24 --anycast-ip 12:10.0.12.1/24 --anycast-ipv6
11:11::1/127 --anycast-ipv6 12:12::1/127 --vrf ten1vrf1
     --suppress-arp 11:false –suppress-nd 12:false

efa tenant epg create --name ten2epg1 --tenant ten2 --port
10.20.246.15[0/11],10.20.246.16[0/11]
     --switchport-mode trunk --ctag-range 21-22
     --anycast-ip 21:10.0.21.1/24 --anycast-ipv6 21:21::1/127 --anycast-ip
22:10.0.22.1/24 --anycast-ipv6 22:22::1/127 --vrf ten2vrf1
     --suppress-arp 21:false –suppress-nd 22:false

efa tenant epg show –detail

================================================================================================
=============================================================
Name          : ten1epg1
Tenant        : ten1
Type     : extension
State         :
Description   :
Ports         : 10.20.246.15[0/1]
              : 10.20.246.16[0/1]
POs           :
Port Property : SwitchPort Mode     : trunk
```

```
              : Native Vlan Tagging : false
NW Policy     : Ctag Range          : 11-12
              : VRF            : ten1vrf1
              : L3Vni               : 8192

+------+-------------------------+-------+---------+--------------+--------------+--------
+----------------------+---------+----------------+--------------+-------------
+------------------+
| Ctag |          Ctag           | L2Vni | BD Name | Anycast IPv4 | Anycast IPv6 |
Suppress|        Local IP      | IPv6 ND |    IPv6 ND     |   IPv6 ND    |  Dev State  |
App State  |
|      |       Description       |       |         |              |              | ARP/ND
|[Device-IP->Local-IP] |   Mtu   | Managed Config | Other Config |
|             |
+------+-------------------------+-------+---------+--------------+--------------
+-------------+----------------------+---------+----------------+--------------
+-------------+-------------+
|  11  | Tenant L3 Extended VLAN |  11   |         | 10.0.11.1/24 |  11::1/127   |
F/T        |                      |         |     false      |    false     | provisioned
| cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+-------------+----------------------+---------+----------------+--------------
+-------------+-------------+
|  12  | Tenant L3 Extended VLAN |  12   |         | 10.0.12.1/24 |  12::1/127   |
T/F        |                      |         |     false      |    false     | provisioned
| cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------
+-------------+----------------------+---------+----------------+--------------
+-------------+-------------+
================================================================================================
=============================================================================
================================================================================================
=============================================================================
Name          : ten2epg1
Tenant        : ten2
Type     : extension
State         :
Description   :
Ports         : 10.20.246.15[0/11]
              : 10.20.246.16[0/11]
POs           :
Port Property : SwitchPort Mode     : trunk
              : Native Vlan Tagging : false
NW Policy     : Ctag Range          : 21-22
              : VRF            : ten2vrf1
              : L3Vni               : 8191
+------+-----------------------+-------+--------------+--------------+--------------
+--------+----------------------+---------+----------------+--------------+-------------
+----------------+
| Ctag |          Ctag         | L2Vni | BD Name      | Anycast IPv4 | Anycast IPv6 |
Suppress|        Local IP      | IPv6 ND |    IPv6 ND     |   IPv6 ND    |  Dev State  |
App State|
|      |       Description     |       |              |              |              |
ARP/ND |[Device-IP->Local-IP] |   Mtu   | Managed Config | Other Config |
|           |
+------+-----------------------+-------+--------------+--------------+--------------
+-------------+----------------------+---------+----------------+--------------
+-------------+-----------+
|  21  | Tenant L3 Extended BD | 4097  | Auto-BD-4097 | 10.0.21.1/24 |  21::1/127   |
F/T        |                       |         |     false      |    false     |
provisioned |cfg-in-sync|
+------+-----------------------+-------+--------------+--------------+--------------
+-------------+----------------------+---------+----------------+--------------
+-------------+-----------+
```

```
|  22  | Tenant L3 Extended BD | 4098  | Auto-BD-4098 | 10.0.22.1/24 |  22::1/127   |
T/F         |                       |         |     false      |    false     |
provisioned |cfg-in-sync|
+------+----------------------+-------+--------------+--------------+--------------
+-------------+----------------------+---------+----------------+--------------
+-------------+-----------+
================================================================================================
=================================================================
```

Configure Local IP for Endpoint Group

You can configure local IP address for an endpoint group.

About This Task

Follow this procedure to add or delete local IP address configurations during the following operations:

- Creating endpoint groups (EPGs)
- Adding or deleting CTAG ranges
- Adding or deleting VRFs

The Local IP address is configured on the VE interface that is assigned to a tenant network. You can select different local IP addresses for each device in a tenant network.

Procedure

1. To configure local IP when you create an EPG, run the following commands:

```
efa tenant epg create --name ten1epg1 --tenant tenant1 --vrf red
--switchport-mode trunk --ctag-range 11 --anycast-ip 11:10.10.11.1/24
--port 10.24.80.150[0/1],10.24.80.151[0/1]
--local-ip 11,10.24.80.150:11.22.33.41/24 --local-ip 11,10.24.80.151:11.22.34.41/24

efa tenant epg show
Name: ten1epg1
Tenant: tenant1
Description:
Type: extension
Ports : 10.24.80.151[0/1]
      : 10.24.80.150[0/1]
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy: ctag-range       :11
              : vrf                 : red
              : vrf-State           : vrf-device-created
              : vrf-Device-State    : provisioned
              : vrf-App-State       : cfg-refreshed
              : l3-vni              : 8190
Network Property [Flags : * - Native Vlan]

| Ctag | L2-Vni | Anycast-IPv4 | Anycast-IPv6 | BD-name |    Local IP (Device-IP-
>Local-IP)   |  Dev-state  |   App-state   |
+------+--------+--------------+--------------+---------
+-------------------------------------+-------------+---------------+
| 11   | 11     |10.10.11.1/24 |
|         |  10.24.80.151->11.22.34.41/24       | provisioned | cfg-refreshed |
|      |        |              |
|         |                                     |             |               |
|      |        |              |
|         |  10.24.80.150->11.22.33.41/24       |             |               |
```

2. To delete local IP when you update an EPG, run the following commands:

```
efa tenant epg update --name epgv20 --tenant tenant1 --operation
local-ip-delete --local-ip 11,10.24.80.150:11.22.33.41/24

efa tenant epg show
Name          : epgv20
Tenant        : t3
Description   :
Type          : l3-hand-off
Ports         : 10.20.50.209[0/27]
POs           : posv9
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 201-202
              : vrf                 : vrfv20
              : l3-vni              : 5110

Network Property [Flags : * - Native Vlan]
+------+--------+--------------+--------------+---------
+-------------------------------------+-------------+---------------+
| Ctag | L2-Vni | Anycast-IPv4 | Anycast-IPv6 | BD-name |    Local IP (Device-IP-
>Local-IP)    |  Dev-state  |   App-state   |
+------+--------+--------------+--------------+---------
+-------------------------------------+-------------+---------------+
| 201  | 201    |              |              |         |  10.20.50.209-
>44.4.4.5/24           | provisioned | cfg-in-sync   |
|      |        |              |              |         |
4444:44::5/120        |             |               |
|      |        |              |              |         |  10.20.50.208-
>44.4.4.4/24           |             |               |
|      |        |              |              |         |
4444:44::4/120        |             |               |
+------+--------+--------------+--------------+---------
+-------------------------------------+-------------+---------------+
| 202  | 202    |              |              |         |  10.20.50.209-
>44.4.5.5/24           | provisioned | cg-in-sync    |
|      |        |              |              |         |
4444:45::5/120        |             |               |
+------+--------+--------------+--------------+---------
+-------------------------------------+-------------+---------------+

For 'unstable' entities, run 'efa tenant po/vrf show' for details
================================================================================================
=========================================
```

3. To add local IP when you update an EPG, run the following commands:

```
efa tenant epg update --name ten1epg1 --tenant tenant1
--operation local-ip-add --local-ip 11,10.24.80.150:11.22.33.41/24

efa tenant epg show
Name: ten1epg1
Tenant: tenant1
Description:
Type: extension
Ports : 10.24.80.151[0/1]
      : 10.24.80.150[0/1]
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy: ctag-range       :11
              : vrf                 : red
              : vrf-State           : vrf-device-created
              : vrf-Device-State    : provisioned
              : vrf-App-State       : cfg-refreshed
              : l3-vni              : 8190
```

```
Network Property [Flags : * - Native Vlan]
+------+--------+-------------+-------------+---------
+-----------------------------------------+-------------+-------------+
| Ctag | L2-Vni | Anycast-IPv4 | Anycast-
IPv6 | BD-name |       Local IP (Device-IP->Local-IP)     |  Dev-state  |  App-state  |
+------+--------+-------------+-------------+---------
+-----------------------------------------+-------------+-------------+
| 11   | 11     |             |
|         |  10.24.80.151->11.22.34.41/24            | provisioned | cfg-in-sync |
|      |        |             |
|         |                                          |             |             |
|      |        |             |
|         |  10.24.80.150->11.22.33.41/24            |             |             |
```

EPG: Network Property: IP MTU

You can configure the maximum transmission unit (MTU) when you create or update an endpoint group.

About This Task

You use the `--ip-mtu` parameter (in the format `ctag:value`) to configure the MTU for the tenant network. This value is then configured on the interface VE on the SLX device.

Procedure

1. To configure MTU when you create an endpoint group, run the **`efa tenant epg create`** command.

   ```
   $ efa tenant epg create --name ten1epg1 --tenant ten1 --port 10.20.246.17[0/1],
   10.20.246.18[0/1] --switchport-mode trunk --ctag-range 11-12 --anycast-
   ip11:10.0.11.1/24
   --anycast-ip12:10.0.12.1/24 --anycast-ipv6 11:11::1/127 --anycast-ipv6 12:12::1/127
   --vrf ten1vrf1 --ip-mtu 11:7900 --ip-mtu 12:8900
   ```

   This example creates an endpoint group with MTU values for Ctag 11 and Ctag 12.
2. To configure MTU for an existing endpoint group, run the **`efa tenant epg update`** command during `vrf-add` or `ctag-range-add` operations.

   ```
   $ efa tenant epg update --name ten1epg1 --tenant ten1 --operation ctag-range-add
   --ctag-range 12 --anycast-ip12:10.0.12.1/24 --anycast-ipv6 12:12::1/127 --ip-mtu
   12:6990
   ```

   This example configures the MTU during a `ctag-range-add` operation.
3. To view the configured MTU for an endpoint group, run the **`efa tenant epg show --detail`** command.

IPv6 Anycast Gateway Support

XCO supports the provisioning of IPv6 anycast gateways.

The following commands support provisioning of anycast gateways in XCO.

```
efa tenant vrf create --name vrf1 --tenant tenant1 --rt-type both --rt 1:1

efa tenant epg create --name ten1epg1 --tenant tenant1 --port 10.24.80.134[0/15]
--switchport-mode trunk --ctag-range 1001 --anycast-ip 1001:10.0.1.1/24
--anycast-ipv6 1001:2001:10:0:1::1/64 --vrf vrf1
```

Software BFD Session Support on CEP

You can configure software Bidirectional Forwarding Detection (BFD) sessions on a Cluster Edge Port (CEP) on SLX 9150 and SLX 9250 devices. The EPG Port Property shows the `bfd-software-session` attribute, using which you can choose a software or hardware BFD session.

BFD Session Formation with SRIOV Server

During initial state of BFD session formation with SRIOV (single-root input or output virtualization) server:

- For MCT-1:
  - The nexthop reachability for 10.1.1.3 is via ICL.
  - It forms a software BFD session with 10.1.1.3 via ICL. It also forms a software BFD session with 10.1.1.3.
- For MCT-2:
  - The nexthop reachability for 10.1.1.3 is via CEP port eth 0/1.
  - It forms a hardware BFD session with 10.1.1.3.

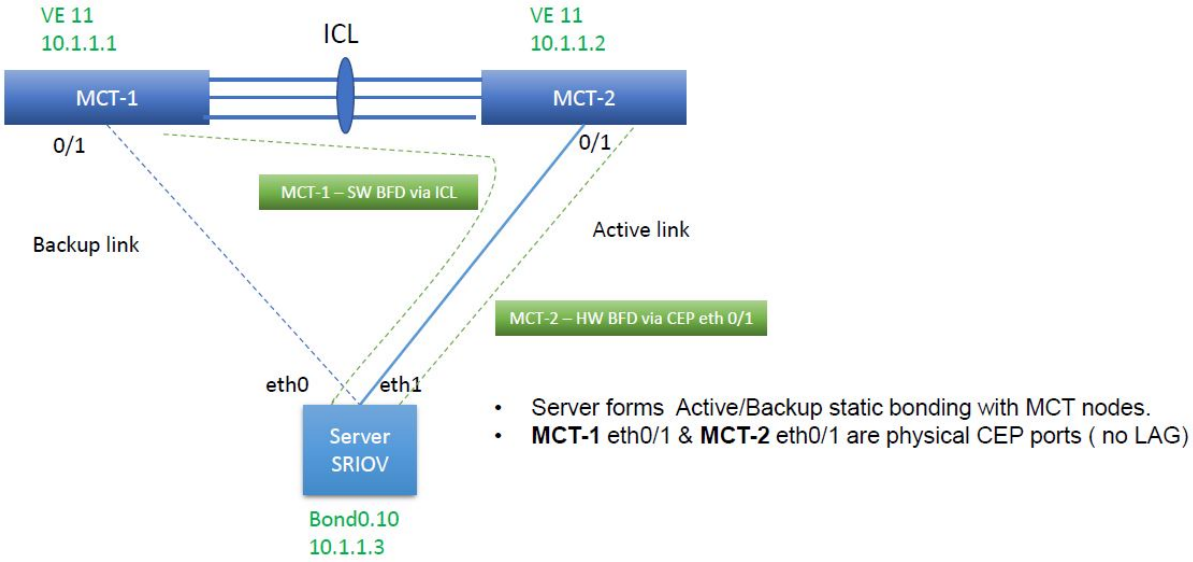

BFD Session Formation with SRIOV Server after Link Failover

During the link failover of BFD session formation with SRIOV (single-root input/output virtualization) server:

- For MCT-1:
  - The Nexthop reachability for 10.1.1.3 changes from ICL to its CEP eth 0/1.
  - The BFD session changes from software to hardware. BFD reachability for 10.1.1.3 changes from ICL to its CEP eth 0/1.
- For MCT-2:
- ◦ The Nexthop reachability for 10.1.1.3 changes from CEP eth 0/1 to ICL.
  - The BFD session changes from hardware to software BFD.

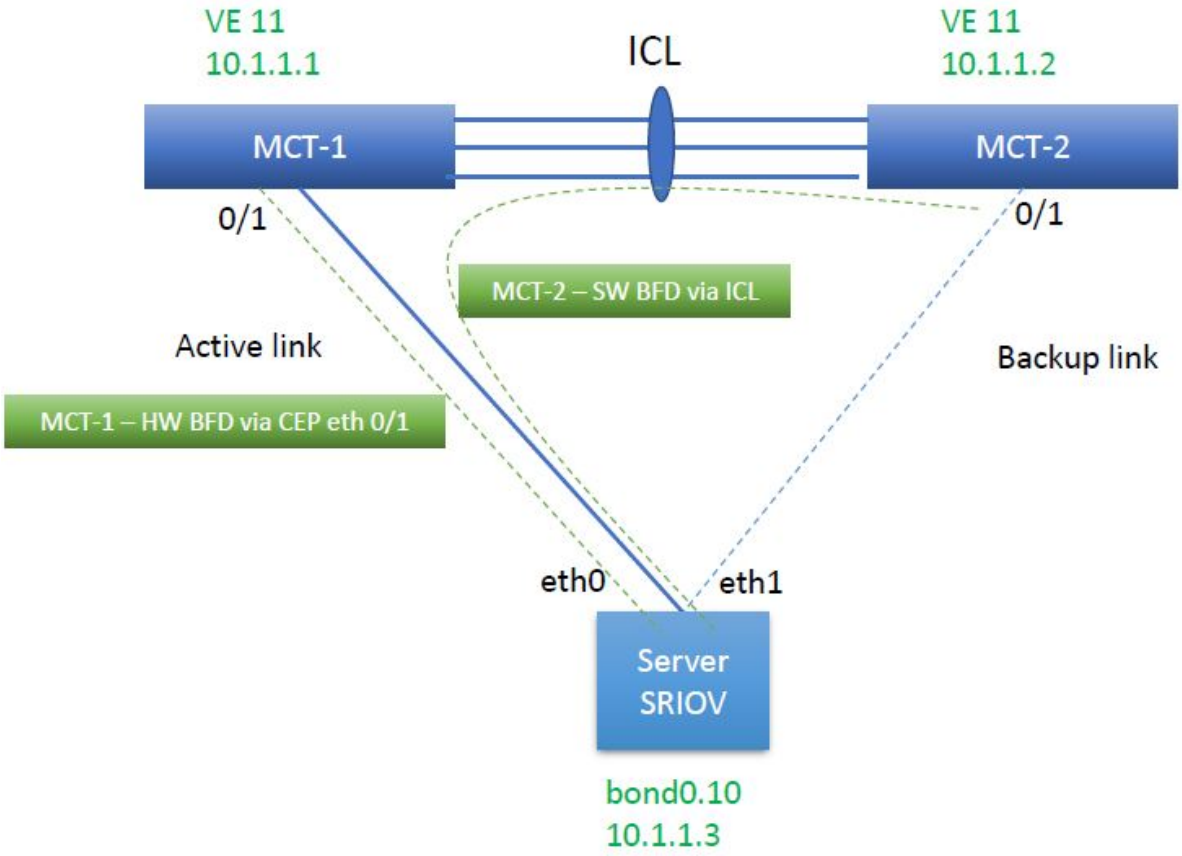

BFD Session Formation with SRIOV Server

- Limitation in SLX 9250 and 8720

  Transition between software and hardware based BFD is not supported. Therefore, during session formation with SRIOV, BFD does not come up during link failover.

  Note
  - Introduce a CLI knob to change the BFD session formed over a CEP (Ethernet or port channel) port to software BFD sessions instead of hardware BFD.
  - With the CLI knob, both MCT 1 and MCT 2 can form a software BFD sessions with SRIOV server.
  - During link failover, it is SW-SW BFD session transition instead of HW-SW BFD session transition.

- SLX Configuration

```
interface Ethernet 0/1
 cluster-track
 bfd-software-session
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
11
 no switchport trunk tag native
vlan
 no shutdown
!
```

```
interface Port-channel 1
 cluster-track
 bfd-software-session
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
11
 no switchport trunk tag native
vlan
 no shutdown
!
```

cep-bfd-session-type Automation on EPG (Endpoint Group) Port Property

XCO automates the `cep-bfd-session-type` on the CEP interfaces based on the logic with no additional input from the users.

SLX Hardware Type	XCO: Fabric Links (Leaf to Spine)	XCO: Extension EPG		XCO: L3-Handoff EPG
CEP SRIOV	CEP Non-SRIOV	CEP		
SLX 9250	Hardware	Software	Software	Hardware
SLX 8720	Hardware	Software	Software	Hardware
SLX 9740 and other SLX versions	Hardware	Hardware	Hardware	Hardware

```
(efa:root)root-2:-# efa fabric show
Fabric Name: default, Fabric Description: Default Fabric, Fabric Stage: 3, Fabric Type:
clos, Fabric Status: created
+--------+-----+-----+-----+------+-------+------+-----------+--------+-----+----+
| IP     | POD | HOST| ASN | ROLE | DEVICE| APP  | CONFIG GEN| PENDING| VTLB| LB |
| ADDRESS|     | NAME|     |      | STATE | STATE| REASON    | CONFIGS| ID  | ID |
+--------+-----+-----+-----+------+-------+------+-----------+--------+-----+----+

Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric Status:
settings-updated

Updated Fabric Settings: BGP-LL

+-------------+---+------+-------+------+--------------------+------------+------------
+---------------+---------+---+
| IP ADDRESS  |POD| HOST | ASN   | ROLE | DEVICE STATE       | APP STATE  |
CONFIG GEN | PENDING       | VTLB ID | LB|
|             |   | Name |       |      |                    |            |
| REASON        | CONFIGS | ID|
+-------------+---+------+-------+------+--------------------+------------+------------
+---------------+---------+---+
| 10.20.246.1 |   | SLX-1| 64512 | Spine| provisioned        | cfg in-sync|
NA         | NA            | NA      | 1 |
| 10.20.246.7 |   | SLX  | 65000 | Leaf | provisioning failed | cfg ready
| IA,IU,MD,DA| SYSP-C,MCT-C  |         |   |
|             |   |      |       |      |                    |            |
| MCT-PA,BGP-C  |         |   |
|             |   |      |       |      |                    |            |
```

```
| INTIP-C,EVPN-C|          |   |
|             |   |       |       |      |                      |            |
| O-C           | 2       | 1 |
| 10.20.246.8 |   | slx-8| 65000 | Leaf | provisioned          | cfg in-sync|
NA          | NA            | 2       | 1 |
+-------------+---+------+-------+------+----------------------+------------+------------
+---------------+---------+---+
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5 Password , BFD-RX -
Bfd Rx Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd multiplier,
BFD-ENABLE - Enable Bfd, BGP-MULTIHOP - BGP ebgp multihop, P2PLR - Point-to-Point
Link Range, MCTLR - MCT Link Range, LOIP - Loopback IP Range


CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn
Delete/Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS
- System Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit,
POU - Port Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface
IP, BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete
(efa:root)root-2:-# efa tenant show
+------+---------+-------+-------+-------+------+--------+----------------------+
| Name |   Type  | VLAN  | L2VNI | L3VNI | VRF  | Enable |         Ports        |
|      |         | Range | Range | Range | Count| BD     |                      |
+------+---------+-------+-------+-------+------+--------+----------------------+
| ten1 | Private | 11-20 |       |       | 10   | false  | 10.20.246.6[0/1-10]| |
|      |         |       |       |       |      |        | 10.20.246.5[0/1-10]| |
+------+---------+-------+-------+-------+------+--------+----------------------+

efa tenant vrf create --name ten1vrf1 --tenant ten1
```

EPG Create

Run the following command to create a `cep-bfd-session-type` automation on EPG port property:

```
efa tenant epg create --name ten1epg1 --tenant ten1
   --port 10.20.246.5[0/1],10.20.246.6[0/1]
   --switchport-mode trunk
   --vrf ten1vrf1 --ctag-range 11
   --anycast-ip 11:20.0.11.1/24
   --local-ip 11,10.20.246.5:10.1.1.1/24 --local-ip 11,10.20.246.6:10.1.1.2/24
```

Example:

```
(efa:root)root@node-2:-# efa tenant epg show  --detail
================================================================
Name          : ten1epg1
Tenant        : ten1
Type          : extension
State         :
Description   :

Ports         : 10.20.246.5[0/1]
              : 10.20.246.6[0/1]
```

```
POs           :
Port Property : SwithchPort Mode    : trunk
              : Native Vlan Tagging : false
              : BFD Session Type    : Auto
NW Policy     : Ctag Range          : 11
              : VRF                 : ten1vrf1
              : L3Vni               : 8192
+------+---------------+-------+-----+-------------+--------+--------------------------
+------+---------------+-------------+-------------+------------+
| Ctag |     Ctag      | L2Vni | BD  | Anycast IPv4 | Anycast|         Local
IP          |IPv6  |     IPv6 ND    |    IPv6 ND   |  Dev State  |  App State |
|      |  Description  |       | Name|              | IPv6   |  [Device-IP->Local-
IP]   |ND Mtu| Managed Config | Other Config |             |            |
+------+---------------+-------+-----+-------------+--------+--------------------------
+------+---------------+-------------+-------------+------------+
|  11  | Tenant L3     |  11   |     | 20/0.11.1/24 |        | 10.20.246.5->10.1.1.1/24
|      |      false    |     false   | provisioned | cfg-in-sync|
|      | Extended VLAN |       |     |              |        | 10.20.246.6->10.1.1.1/24
|      |               |             |             |            |
+------+---------------+-------+-----+-------------+--------+--------------------------
+------+---------------+-------------+-------------+------------+
Network Property [Flags : * - Native Vlan]
+------+----------------+-------------+----------------+--------------------
+-------------+
| CTAG | IPv6 ND Prefix | No Advertise
| Valid Lifetime | Preferred Lifetime | Config Type |
+------+----------------+-------------+----------------+--------------------
+-------------+
IPv6 ND Prefix Flags

For 'unstable' entities, run 'efa tenant po/vrf show' for details
================================================================================
```

VRF Update

Run the following command to update a tenant VRF on the static route BFD:

```
efa tenant vrf update --name ten1vrf1 --tenant ten1
        --operation static-route-bfd-add
        --ipv4-static-route-bfd 10.20.246.5,10.1.1.3,10.1.1.1
        --ipv4-static-route-bfd 10.20.246.6,10.1.1.3,10.1.1.2
```

Example

```
(efa:root)root@node-2:-# efa tenant vrf show  --detail
================================================================
Name                    : ten1vrf1
Tenant                  : ten1
Type                    : extension
Routing Type            : distributed
Centralized Routers     :
Redistribute            : connected
Max Path                : 8
Local ASN               :
L3VNI                   : 8192
EVPN IRB BD             : 4096
EVPN IRB VE             : 8192
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            :
Static Rout BFD         : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                        : 10.20.246.6->10.1.1.3,10.1.1.2
                        : 10.20.246.5->10.1.1.3,10.1.1.1
State                   : vrf-device-created
Dev State               : provisioned
App State               : cfg-in-sync

================================================================================
```

Switch Config

```
Rack1-Device1(config)# do show
running-config vlan 11
vlan 11
 router interface Ve 11
 suppress-arp
 description Tenant L3 Extended
VLAN
!
Rack1-Device1(config)# do show
running config interface Ve 11
interface Ve 11
 vrf forwarding ten1vrf1
 ip anycast address 20.0.11.1/24
 ip address 10.1.1.1/24
 no shutdown
!
Rack1-Device1(config)# do show
running config interface Ethernet
0/1
interface Ethernet 0/1
 cluster-track
 bfd software session
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
11
 no switchport trunk tag native-
vlan
 no shutdown
!
Rack1-Device1(config)# do show
running config vrf ten1vrf1
address family
ipv4 unicast
vrf ten1vrf1
 address family ipv4 unicast
  route target export 101:101 evpn
  route target import 101:101 evpn
  ip route static bfd 10.1.1.3
10.1.1.1
 !
!
Rack1-Device1(config)#
```

```
Rack1-Device2(config)# do show
running config vlan 11
vlan 11
 router interface Ve 11
 suppress arp
 description Tenant L3 Extended
VLAN
!
Rack1Device2(config)# do show
running config in Ve 11
interface Ve 11
 vrf forwarding ten1vrf1
 ip anycast address 20.0.11.1/24
 ip address 10.1.1.2/24
 no shutdown
!
Rack1Device2(config)# do show
running config int eth 0/1
interface Ethernet 0/1
 cluster track
 bfd software session
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add
11
 no switchport trunk tag native-
vlan
 no shutdown
!
Rack1-Device2(config)# do show
running config vrf ten1vrf1
address family
ipv4 unicast
vrf ten1vrf1
 address family ipv4 unicast
  route target export 101:101 evpn
  route target import 101:101 evpn
  ip route static bfd 10.1.1.3
10.1.1.2
 !
!
Rack1-Device2(config)#
```

cep-bfd-session-type on EPG Port Property

- The `cep-bfd-session-type` enables you to provide `cep bfd session type` per EPG which gets configured on all the ethernet and single homed port channel interfaces defined in the EPG.

- The default value of `cep bfd session type` is set to “auto”. XCO automatically derives the appropriate `cep bfd session type` value based on the use case (extension or l3 hand off) and endpoint type.
- You can provide the `cep bfd session type` configuration when you create or update an EPG and add port group.

Operation	Command

Create EPG	`efa tenant epg create --name < epg-name> --tenant <tenant-name>`
	`--port <port-list> po <po-list>`
	`--switchport-mode <access | trunk>` **`--cep-bfd-session-type {auto | software | hardware}`**

Update EPG	`efa tenant epg update --name < epg-name> tenant <tenant-name>`
	`--operation port group add`
	`--port <port-list> --po <po-list>`
	`switchport-mode < access|trunk >` **`--cep-bfd-session-type {auto | software | hardware}`**

CEP SRIOV and Non-SRIOV	Upgrade Handling
• XCO cannot distinguish between the SRIOV and non-SRIOV connections. Hence both the CEP SRIOV and CEP non-SRIOV phy or port channel are treated in same manner. • To use a “hardware” BFD sessions for the CEP non-SRIOV connections, create an EPG containing all the non-SRIOV CEP with `cep bfd session type=hardware`.	During upgrade from EFA 2.5.5 and onwards, all the CEP ports (on SLX 9250 and SLX 8720 platforms) used in the “extension” EPG must have `cep bfd session type` as `software`. You must perform an explicit DRC to reconcile the XCO configuration to synchronize with the SLX.

Co-existence of centralize and distributed routing on a CEP

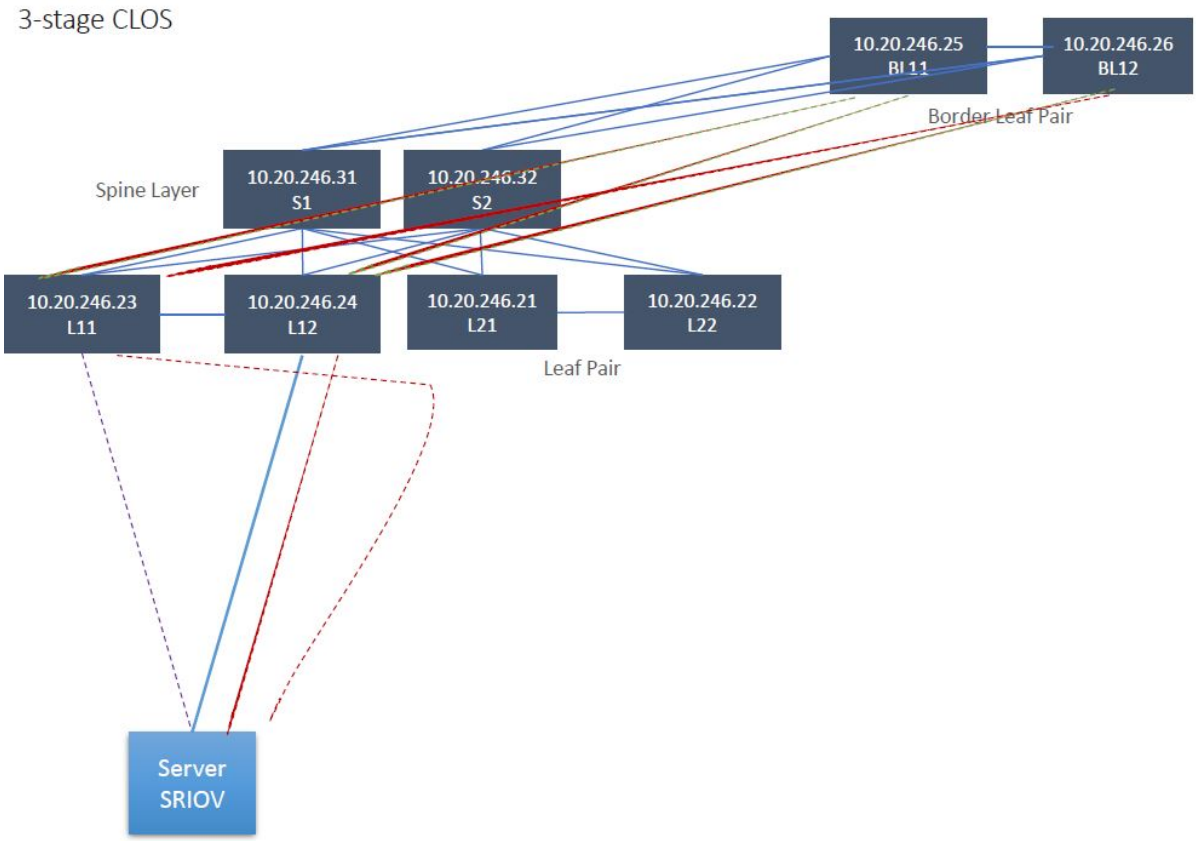

Provision a BGP Peer

You can configure a BGP peer.

About This Task

Complete the following tasks to configure a BGP peer in your XCO fabric:

Procedure

1. Create BGP Static Peer on page 293
2. Create BGP Dynamic Peer on page 301
3. Getting the Operational State of the BGP Peers on page 308
4. Configure Route Map Attribute on page 310
5. Configure remove-private-as on BGP Peer on page 313
6. Configure default-originate to Advertise Default Route on BGP Peer on page 316
7. Configure Backup Routing Neighbors on BGP Peer on page 320
8. Configure Send-Community on Tenant BGP Peer on page 321
9. Configure Out-of-band for a Tenant BGP Peer or Peer Group on page 324

Create BGP Static Peer

You can configure a BGP static peer when you create or update a BGP peer.

About This Task

Follow this procedure to configure BGP static peer.

Procedure

1. To create a BGP static peer when you create a BGP peer, run the following command:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
   --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
   --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
   --ipv4-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,
           all|both|extended|large|standard|large-and-standard|large-and-extended>
   --ipv6-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,
           all|both|extended|large|standard|large-and-standard|large-and-extended>
```

   The following example creates a BGP static peer:

```
efa tenant service bgp peer create --name ten1bgppeer1 --tenant ten1
   --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
   --ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
   --ipv4-uc-nbr-send-community 10.20.246.15,ten1vrf1:10.20.30.40,all
   --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
   --ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
   --ipv4-uc-nbr-send-community 10.20.246.16,ten1vrf1:10.20.30.40,both
```

2. To view BGP static peer, run the following command:

```
efa tenant service bgp peer show --
detail
==================================
====
Name              : ten1bgppeer1
Tenant            : ten1
State             : bs-state-created
Description       :

Static Peer
-----------
  Device IP         : 10.20.246.15
  VRF               : ten1vrf1
  AFI               : ipv4
  SAFI              : unicast
  Remote IP         : 10.20.30.40
  Remote ASN        : 50000
  Next Hop Self     : false
  Update Source IP  :
  BFD Enabled       : true
  BFD Interval      : 0
  BFD Rx            : 0
  BFD Multiplier    : 0
  Default Originate:
  Default Originate
  Route Map         :
  Send Community    : all
  Dev State         : provisioned
  App State         : cfg-in-sync

  Device IP         : 10.20.246.15
  VRF               : ten1vrf1
  AFI               : ipv4
  SAFI              : unicast
  Remote IP         : 10.20.30.50
  Remote ASN        : 50000
  Next Hop Self     : false
  Update Source IP  :
  BFD Enabled       : true
  BFD Interval      : 0
  BFD Rx            : 0
  BFD Multiplier    : 0
  Default Originate: false
  Default Originate
  Route Map         :
  Send Community    : standard

  Dev State         : provisioned
  App State         : cfg-in-sync
```

```
Device IP         : 10.20.246.16
  VRF               : ten1vrf1
  AFI               : ipv4
  SAFI              : unicast
  Remote IP         : 10.20.30.40
  Remote ASN        : 50000
  Next Hop Self     : false
  Update Source IP  :
  BFD Enabled       : true
  BFD Interval      : 0
  BFD Rx            : 0
  BFD Multiplier    : 0
  Default Originate: false
  Default Originate
  Route Map         :
  Send Community    : both
  Dev State         : provisioned
  App State         : cfg-in-sync

  Device IP         : 10.20.246.16
  VRF               : ten1vrf1
  AFI               : ipv4
  SAFI              : unicast
  Remote IP         : 10.20.30.50
  Remote ASN        : 50000
  Next Hop Self     : false
  Update Source IP  :
  BFD Enabled       : true
  BFD Interval      : 0
  BFD Rx            : 0
  BFD Multiplier    : 0
  Default Originate: false
  Default Originate
  Route Map         :
  Send Community    : extended
  Dev State         : provisioned
  App State         : cfg-in-sync
Dynamic Peer
-----------
        0 Records



        0 Records
==================================
```

3. To add a BGP static peer when you update a BGP peer, run the following command:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
  --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
  --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-interval,bfd-
rx,bfdmult>
  --ipv4-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,all|both|extended|large|
standard|large-and-standard|large-and-extended>
  --ipv6-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,all|both|extended|large|
standard|large-and-standard|large-and-extended>
```

The following example adds a BGP static peer:

```
# efa tenant service bgp peer update --name bgpservice1 --tenant
tenant1 --operation peer-add --ipv6-uc-nbr 10.24.80.134,red:10::40,5000 --ipv6-
uc-nbr-bfd 10.24.80.134,red:10::40,true,100,200,5 --ipv6-uc-nbr-update-source-ip
10.24.80.134,red:10::40,11::22 --ipv6-uc-nbr-next-hop-self 10.24.80.134,red:10::40,true
efa tenant service bgp peer show
================================================================================================
====
Name : bgpservice1
Tenant : tenant1
State : bs-state-created
+-------------+---+-----+--------+-------+-------+-------+--------+----+----------
+-----------+-----------+
|Device IP    |VRF| AFI | SAFI   | REMOTE| REMOTE|BFD    |BFD     |BFD |BFD       |Dev-
state  |  App-state|
|             |   |     |        | IP    | ASN   |Enabled|Interval|Rx  |
Multiplier|           |           |
+-------------+---+-----+--------+-------+-------+-------+--------+----+----------
+-----------+-----------+
|10.24.80.134 |red| ipv6| unicast| 10::40| 5000  |false  |100     |200 |5         |
provisioned|cfg-in-sync|
+-------------+---+-----+--------+-------+-------+-------+--------+----+----------
+-----------+-----------+
================================================================================================
====
```

4. Verify the switch configuration on SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.123/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 send-
community all
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 send-
community standard
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.176/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 send-
community both
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 send-
community extended
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

5. To delete a BGP static peer when you update a BGP peer, run the following command:

```
# efa tenant service bgp peer update --name <peer-name> --tenant <tenant-name>
--operation peer-delete --ipv4-unicast-neighbor <switch-ip,vrf-name:ipv4-neighbor> --
ipv6-unicast-neighbor <switch-ip,vrf-name:ipv4-neighbor >
```

The following example deletes a BGP peer:

```
# efa tenant service bgp peer update --name bgpservice1 --tenant tenant1 --operation
peer-delete --ipv4-unicast-neighbor 10.24.80.134,red:10.20.30.40

# efa tenant service bgp show
================================================================================================
==
Name : bgpservice1
Tenant : tenant1
State : bs-state-created
+------------+----+----+-------+-------+-------+--------+---------+----+----------
+------------+-----------+
|Device IP   |VRF |AFI |SAFI   |REMOTE |REMOTE |BFD     |BFD      |BFD |BFD       |Dev-
```

```
state   |App-state  |
|            |    |    |       |IP     |ASN    |Enabled |Interval |Rx  |
Multiplier|            |           |
+------------+----+----+-------+-------+-------+--------+---------+----+----------
+------------+-----------+
|10.24.80.134|red |ipv6|unicast|10::40 |5000   |false   |0        |0   |0         |
provisioned |cfg-in-sync|
+------------+----+----+-------+-------+-------+--------+---------+----+----------
+------------+-----------+
================================================================================================
==
```

Add Path on Tenant BGP Peer

You can configure additional paths (for both IPv4 and IPv6) when you create or update a BGP peer.

About This Task

Follow this procedure to add paths on tenant BGP peer.

Procedure

1. To configure an additional path when you create a BGP peer, run the following command:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
        --ipv6-uc-nbr-add-path-capability <device-ip,vrf-name:neighbor-ip,
               {send | receive | both}>
        --ipv6-uc-nbr-add-path-advertise-all <device-ip,vrf-name:neighbor-ip,
               {true | false}>
        --ipv6-uc-nbr-add-path-advertise-group-best <device-ip,vrf-name:neighbor-ip,
               {true | false}>
        --ipv6-uc-nbr-add-path-advertise-best <device-ip,vrf-name:neighbor-ip,
               2-16>
```

2. To configure an additional path when you update a BGP peer, run the following command:

```
efa tenant service bgp peer update --name <bgp-peer-name> --tenant <tenant-name>
        --ipv4-uc-nbr-add-path-capability <device-ip,vrf-name:neighbor-ip,
               {send | receive | both}>
        --ipv4-uc-nbr-add-path-advertise-all <device-ip,vrf-name:neighbor-ip,
               {true | false}>
        --ipv4-uc-nbr-add-path-advertise-group-best <device-ip,vrf-name:neighbor-ip,
               {true | false}>
        --ipv4-uc-nbr-add-path-advertise-best <device-ip,vrf-name:neighbor-ip,
               2-16>
```

3. Verify the configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.123/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  additional-paths select all
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 additional-
paths send receive
  neighbor 10.20.30.40 additional-
paths advertise best 10 group-best
all
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 additional-
paths send
  neighbor 10.20.30.50 additional-
paths advertise best 8 group-best
all
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.176/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  additional-paths select all
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 additional-
paths send receive
  neighbor 10.20.30.40 additional-
paths advertise best 5
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 additional-
paths receive
  neighbor 10.20.30.50 additional-
paths advertise best 4
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Example

The following is an example output for adding additional paths when you create or update a BGP peer:

```
efa tenant service bgp peer create --name ten1bgppeer1 --tenant ten1
   --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
   --ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
   --ipv4-uc-nbr-add-path-capability 10.20.246.15,ten1vrf1:10.20.30.40,both
   --ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,ten1vrf1:10.20.30.40,true
```

```
   --ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,ten1vrf1:10.20.30.40,true
   --ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,ten1vrf1:10.20.30.40,10

   --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
   --ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
   --ipv4-uc-nbr-add-path-capability 10.20.246.16,ten1vrf1:10.20.30.40,both
   --ipv4-uc-nbr-add-path-advertise-all 10.20.246.16,ten1vrf1:10.20.30.40,false
   --ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.16,ten1vrf1:10.20.30.40,false
   --ipv4-uc-nbr-add-path-advertise-best 10.20.246.16,ten1vrf1:10.20.30.40,5


efa tenant service bgp peer update --name ten1bgppeer1 --tenant ten1
   --operation peer-add
   --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
   --ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
   --ipv4-uc-nbr-add-path-capability 10.20.246.15,ten1vrf1:10.20.30.50,send
   --ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,ten1vrf1:10.20.30.50,true
   --ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,ten1vrf1:10.20.30.50,true
   --ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,ten1vrf1:10.20.30.50,8


   --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
   --ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.50,true
   --ipv4-uc-nbr-add-path-capability 10.20.246.16,ten1vrf1:10.20.30.50,receive
   --ipv4-uc-nbr-add-path-advertise-all 10.20.246.16,ten1vrf1:10.20.30.50,false
```

```
   --ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.16,ten1vrf1:10.20.30.50,false
   --ipv4-uc-nbr-add-path-advertise-best 10.20.246.16,ten1vrf1:10.20.30.50,4
```

```
efa tenant service bgp peer show --detail
==================================================
Name              : ten1bgppeer1
Tenant            : ten1
State             : bs-state-created
Description       :

Static Peer
-----------
 Device IP           : 10.20.246.15
 VRF                 : ten1vrf1
 AFI                 : ipv4
 SAFI                : unicast
 Remote IP           : 10.20.30.40
 Remote ASN          : 50000
 Next Hop Self       : false
 Update Source IP    :
 BFD Enabled         : true
 BFD Interval        : 0
 BFD Rx              : 0
 BFD Multiplier      : 0
 Default Originate   :
 Default Originate
 Route Map           :
 Add Path Capability : Send,
Receive
 Add Path Advertise  : All, Group
Best, Best 10
 Dev State           : provisioned
 App State           : cfg-in-sync

 Device IP           : 10.20.246.15
 VRF                 : ten1vrf1
 AFI                 : ipv4
 SAFI                : unicast
 Remote IP           : 10.20.30.50
 Remote ASN          : 50000
 Next Hop Self       : false
 Update Source IP    :
 BFD Enabled         : true
 BFD Interval        : 0
 BFD Rx              : 0
 BFD Multiplier      : 0
 Default Originate   : false
 Default Originate
 Route Map           :
Add Path Capability : Send
 Add Path Advertise  : All, Group
Best, Best 8

 Dev State           : provisioned
 App State           : cfg-in-sync
```

```
Device IP        : 10.20.246.16
 VRF             : ten1vrf1
 AFI             : ipv4
 SAFI            : unicast
 Remote IP       : 10.20.30.40
 Remote ASN      : 50000
 Next Hop Self   : false
 Update Source IP :
 BFD Enabled     : true
 BFD Interval    : 0
 BFD Rx          : 0
 BFD Multiplier  : 0
 Default Originate: false
 Default Originate
 Route Map       :
 Add Path Capability : Send,
Receive
 Add Path Advertise  : Best 5
 Dev State       : provisioned
 App State       : cfg-in-sync

 Device IP       : 10.20.246.16
 VRF             : ten1vrf1
 AFI             : ipv4
 SAFI            : unicast
 Remote IP       : 10.20.30.50
 Remote ASN      : 50000
 Next Hop Self   : false
 Update Source IP :
 BFD Enabled     : true
 BFD Interval    : 0
 BFD Rx          : 0
 BFD Multiplier  : 0
 Default Originate: false
 Default Originate
 Route Map       :
 Add Path Capability : Receive
 Add Path Advertise  : Best 4
 Dev State       : provisioned
 App State       : cfg-in-sync
Dynamic Peer
-----------
    0 Records



    0 Records
==============================
```

Create BGP Dynamic Peer

You can configure BGP dynamic peer on a tenant BGP peer.

About This Task

Follow this procedure to configure a BGP dynamic peer.

Procedure

1. To create a BGP dynamic peer when you create a tenant BGP peer, run the following command:

   The listen limit is an optional attribute when you create a BGP dynamic peer.

```
# efa tenant service bgp peer create --name <peer-name> --tenant <tenant-name>
    --ipv4-uc-dyn-nbr <switch-ip,vrf-name:listen-range,peer-group-name,listen-limit>
    --ipv6-uc-dyn-nbr <switch-ip,vrf-name:listen-range,peer-group-name,listen-limit>
```

   Example

```
# efa tenant service bgp peer create --name bgpservice1 --tenant tenant1 --operation
peer-add -ipv4-uc-dyn-nbr 10.24.80.134,red:11::22/127,pg1,20
```

2. To view a BGP dynamic peer, run the following command:

```
# efa tenant service bgp peer show
================================================================================
Name : bgpservice1
Tenant : tenant1
State : bs-state-created
+--------------+-----+------+---------+----------------+------+-------+------------
+------------+
| Device IP    | VRF | AFI  | SAFI    | LISTEN RANGE   | Peer |
LISTEN| Dev-state  | App-state  |
|              |     |      |         |                | Group| LIMIT |
|            |
+--------------+-----+------+---------+----------------+------+-------+------------
+------------+
| 10.24.80.134 | red | ipv4 | unicast | 11.22.33.44/30 | pg1  |
10    | provisioned| cfg-in-sync|
+--------------+-----+------+---------+----------------+------+-------+------------
+------------+
| 10.24.80.134 | red | ipv6 | unicast | 11::22/127     | pg1
| 20    | provisioned| cfg-in-sync|
+--------------+-----+------+---------+----------------+------+-------+------------
+------------+
================================================================================
```

3. Verify the switch configuration on SLX device.

```
Rack1-Device1# sh run router bgp
router bgp
local-as 100
neighbor pg1 peer-group
neighbor pg1 remote-as 6000
address-family ipv4 unicast
!
address-family ipv4 unicast vrf red
listen-range 11.22.33.44/30 peer-group pg1 limit 10
!
address-family ipv6 unicast
!
address-family ipv6 unicast vrf red
listen-range 11::22/127 peer-group pg1 limit 20
!
address-family l2vpn evpn
```

```
!
!
```

4. To delete a dynamic peer when you update a BGP peer, run the following command:

```
# efa tenant service bgp peer update --name <peer-name> --tenant <tenant-name>
--operation peer-delete --ipv4-uc-dyn-nbr <switch-ip,vrf-name:listen-range,peer-group-
name> --ipv6-uc-dyn-nbr <switch-ip,vrf-name:listen-range,peer-group-name>
```

Example

```
# efa tenant service bgp peer create --name bgpservice1 --tenant tenant1 --operation
peer-delete -ipv4-uc-dyn-nbr 10.24.80.134,red:11::22/127,pg1

# efa tenant service bgp peer show
================================================================================
Name   : bgpservice1
Tenant : tenant1
State  : bs-state-created
+-------------+-----+-----+--------+---------------+-----------+--------+------------
+-------------+
| Device IP   | VRF | AFI | SAFI   | LISTEN RANGE  | Peer Group| LISTEN
| Dev-state  | App-state  |
|             |     |     |        |               | Group     | LIMIT  |
|            |
+-------------+-----+-----+--------+---------------+-----------+--------+------------
+-------------+
| 10.24.80.134| red | ipv4| unicast| 11.22.33.44/30| pg1       | 10     | provisioned|
cfg-in-sync|
+-------------+-----+-----+--------+---------------+-----------+--------+------------
+-------------+
================================================================================
```

5. To delete a BGP peer, run the following command:

```
efa tenant service bgp peer delete --name <peer-name> --tenant <tenant-name>
```

Example

```
# efa tenant service bgp peer delete -name bgpservice1 -tenant tenant1
```

Configure Listen Limit on BGP Dynamic Peer

XCO enables configuration of global router BGP listen-limit which depends on the router BGP dynamic peer scale requirements. Listen-limit configuration defined under the “global router bgp” context signifies the maximum number of dynamic BGP peers that can be operational across the VRFs in the SLX.

About This Task

Follow this procedure to configure listen limit.

Default value of the global router BGP listen-limit is 100, which you can modify with any value in the range of 1-2400.

Note

- For the SLX version 20.3.4 or lower, the supported listen-limit range is <1-1024>.
- For the SLX version 20.4.1 or higher, the supported listen-limit range is <1-2400>.

Ensure that the listen-limit configuration defined at the “dynamic peer listen-range” level is less than or equal to the “listen-limit” configuration defined under the “global router bgp”.

The maximum number of dynamic BGP peers is limited to the SLX default (100), when provisioned through XCO.

Configure the new fabric setting for each fabric using the `bgp-dynamic-peer-listen-limit` command. The fabric setting is applicable for Clos and small data center fabrics, and for all types of devices (leaf, border-leaf, spine, super-spine).

Configure the `bgp-dynamic-peer-listen-limit` value on all the devices of fabric when you configure the fabric.

You can configure the `bgp-dynamic-peer-listen-limit` value on an already provisioned fabric. For the new value to be effective, run the **`fabric configure`** command, followed by the **`efa fabric setting update`** command. This enables you to configure **`bgp-dynamic-peer-listen-limit`** on the existing pre-2.5.0 deployments.

Note
You can only increase (but not decrease) the value of `bgp-dynamic-peer-listen-limit` on an already provisioned fabric.

Procedure

1. To configure listen limit on BGP dynamic peer, run the following command:

```
efa fabric setting update -name <fabric-name> --bgp-dynamic-peer-listen-limit <1-2400>
```

2. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 listen-limit 200
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.139/32
  network 172.31.254.226/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 listen-limit 200
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.115/32
  network 172.31.254.226/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Force Delete the Associate Dynamic Peers on a Tenant BGP Peer Group

You can force delete a BGP peer-group to delete the associated dynamic peers.

About This Task

Follow this procedure to forcefully delete an associate dynamic peer.

Procedure

1. Run the **`efa tenant service bgp peer-group show`** command.

```
(efa:root)root@node-2:~# efa tenant service bgp peer-group show
================================================================================================
========
Name        : ten1bgppg1
Tenant      : ten1
State       : bgp-pg-created

+--------------+------+-------+---------+---------+--------+-------------------------
+-----------+------------+
|  Device IP   | Peer | Remote| Next Hop| Update  |   BFD  |           BFD           |
Dev State |  App State |
|              | Group|  ASN  |   Self  |Source IP| Enabled|
[Interval,Rx,Multiplier]|           |            |
+--------------+------+-------+---------+---------+--------+-------------------------
+-----------+------------+
| 10.20.246.16 | pg1  | 65002 |  false  |         |  false |                         |
provisioned| cfg-in-sync|
+--------------+------+-------+---------+---------+--------+-------------------------
+-----------+------------+
| 10.20.246.15 | pg1  | 65002 |  false  |         |  false |                         |
provisioned| cfg-in-sync|
+--------------+------+-------+---------+---------+--------+-------------------------
+-----------+------------+
BGP PeerGroup Details
================================================================================================
==========
 (efa:root)root@node-2:~# efa tenant service bgp peer show
================================================================================================
==========
Name        : ten1bgppeer1
Tenant      : ten1
State       : bgp-peer-created

+------+---+---+-----+-------+-------+----------+---------+--------
+-------------------------+-----+-----+
|Device|VRF|AFI| SAFI| Remote| Remote| Next Hop |Update   |   BFD  |
BFD           |Dev  |App  |
|IP    |   |   |     | IP    | ASN   |   Self   |Source IP| Enabled|
[Interval,Rx,Multiplier]|State|State|
+------+---+---+-----+-------+-------+----------+---------+--------
+-------------------------+-----+-----+
Static Peer Details


+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
|  Device-IP   |   VRF    | AFI  |  SAFI  | Listen Range  | Listen| Peer |  Dev State
|  App State  |
|              |          |      |        |               | Limit | Group|
|             |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.15 | ten1vrf1 | ipv4 | unicast| 10.20.30.0/23 |  100  | pg1  |
```

```
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.15 | ten1vrf1 | ipv4 | unicast| 10.20.40.0/23 |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.15 | ten1vrf1 | ipv6 | unicast|   10::/126    |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.15 | ten1vrf1 | ipv6 | unicast|   20::/126    |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.16 | ten1vrf1 | ipv6 | unicast|   10::/126    |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.16 | ten1vrf1 | ipv6 | unicast|   20::/126    |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.16 | ten1vrf1 | ipv4 | unicast| 10.20.30.0/23 |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
| 10.20.246.16 | ten1vrf1 | ipv4 | unicast| 10.20.40.0/23 |  100  | pg1  |
provisioned| cfg-in-sync |
+--------------+----------+------+--------+--------------+-------+------+------------
+-------------+
Dynamic Peer Details
```

```
================================================================================
==========
```

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor 10.20.20.6 remote-as
4200000000
 neighbor 10.20.20.6 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.153/32
  network 172.31.254.238/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  listen-range 10.20.30.0/23 peer-
group pg1 limit 100
  listen-range 10.20.40.0/23 peer-
group pg1 limit 100
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  listen-range 10::/126 peer-group
pg1 limit 100
  listen-range 20::/126 peer-group
pg1 limit 100
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor 10.20.20.7 remote-as
4200000000
 neighbor 10.20.20.7 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.157/32
  network 172.31.254.238/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  listen-range 10.20.30.0/23 peer-
group pg1 limit 100
  listen-range 10.20.40.0/23 peer-
group pg1 limit 100
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  listen-range 10::/126 peer-group
pg1 limit 100
  listen-range 20::/126 peer-group
pg1 limit 100
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

2. Run the **`efa tenant service bgp peer-group delete`** command.
   If the deletion fails, run the **`efa tenant service bgp peer-group delete`** command with `force` option.

```
(efa:root)root@node-2:~# efa tenant service bgp peer-group delete --name ten1bgppg1
--tenant ten1
BgpService deletion Failed:
Error :  PeerGroup pg1 has dynamic Neighbor 10.20.30.0/23 configured on Device
10.20.246.16


(efa:root)root@node-2:~# efa tenant service bgp peer-group delete --name ten1bgppg1
--tenant ten1 -force

Bgp service peer-group delete with "force" option will delete the device configuration
corresponding to the bgp and also deletes the bgp record from the application. Do you
want to proceed (Y/N): Y

BgpService deleted successfully.
```

```
(efa:root)root@node-2:~# efa tenant service bgp peer-group show
--- Time Elapsed: 192.345588ms ---
(efa:root)root@node-2:~# efa tenant service bgp peer show
================================================================================
Name        : ten1bgppeer1
Tenant      : ten1
State       : bgp-peer-created

+-------+----+-----+-----+-------+-------+----------+-----------+---------
+--------------------------+------+------+
| Device| VRF| AFI | SAFI| Remote| Remote|
Next Hop |  Update   |   BFD   |           BFD            | Dev  | App  |
| IP    |    |     |     | IP    | ASN
|   Self   | Source IP | Enabled | [Interval,Rx,Multiplier] | State| State|
+-------+----+-----+-----+-------+-------+----------+-----------+---------
+--------------------------+------+------+
Static Peer Details

+-----------+-----+-----+------+--------+--------+------+------+------+
| Device-IP | VRF | AFI | SAFI | Listen | Listen | Peer | Dev  | App  |
|           |     |     |      | Range  | Limit  | Group| State| State|
+-----------+-----+-----+------+--------+--------+------+------+------+
Dynamic Peer Details
================================================================================
```

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.6 remote-as
4200000000
 neighbor 10.20.20.6 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.153/32
  network 172.31.254.238/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.7 remote-as
4200000000
 neighbor 10.20.20.7 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.157/32
  network 172.31.254.238/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Getting the Operational State of the BGP Peers

You can get an operational state of the BGP peers that belong to the tenant VRF (non-default VRF).

About This Task

Follow this procedure to get an operational state of a BGP peer.

Procedure

1. Run the following command to create a BGP peer on tenant VRF:

```
efa inventory device register --ip 10.20.246.23,10.20.246.24 --username admin --
password password
efa inventory device register --ip 10.20.246.21,10.20.246.22 --username admin --
password password
efa inventory device register --ip 10.20.246.14 --username admin --password password
efa fabric create --name fabric4 --type clos
efa fabric setting update --name fabric4 --vni-auto-map No --backup-routing-enable yes
efa fabric device add --ip 10.20.246.14 --role spine --name fabric4 --username admin
--password password
efa fabric device add --ip 10.20.246.23 --role leaf  --name fabric4 --username admin
--password password
efa fabric device add --ip 10.20.246.24 --role leaf  --name fabric4 --username admin
--password password
efa fabric device add --ip 10.20.246.22 --role border-leaf --name fabric4 --username
admin --password password
efa fabric device add --ip 10.20.246.21 --role border-leaf --name fabric4 --username
admin --password password
efa fabric configure --name fabric4


efa tenant create --name tenant2 --port 10.20.246.23[0/21-24],10.20.246.24[0/21-24] --
vlan-range 100-200 --l2-vni-range 12000-13000  --vrf-count 25 --l3-vni-range 8000-9000

efa tenant po create --name po101 --port 10.20.246.23[0/22],10.20.246.24[0/22] --speed
10Gbps --negotiation active --tenant tenant2

efa tenant vrf create --name vrf101 --tenant tenant2 --rt-type import --rt 101:102
--rt-type export --rt 105:104
efa tenant vrf create --name vrf102 --tenant tenant2 --rt-type import
--rt 104:105 --rt-type export --rt 200:108 --local-asn 34566 --ipv4-
static-route-next-hop 10.20.246.23,50.0.0.0/24,20.0.0.2 --ipv6-static-route-next-hop
10.20.246.23,3001:1::/64,01::2

efa tenant epg create --name epg1 --ctag-range  100-102 --po po101 --port
10.20.246.23[0/23] --switchport-mode trunk --tenant tenant2 --vrf vrf102 --anycast-ip
100:10.10.10.254/24 --anycast-ip 101:10.10.1.254/24 --anycast-ip 102:10.10.2.254/24
--anycast-ipv6 100:3001:10:0:1::1/64  --anycast-ipv6 101:3001:10:0:2::1/64 --anycast-
ipv6 102:3002:10:0:3::1/64 --local-ip 100,10.20.246.23:121.10.1.2/24 --local-ip
101,10.20.246.23:121.10.2.2/24 --local-ipv6 102,10.20.246.23:121:a::2/64 --local-ip
100,10.20.246.24:121.10.1.3/24 --local-ip 101,10.20.246.24:121.10.2.3/24 --local-ipv6
102,10.20.246.24:121:a::3/64

efa tenant service bgp peer-group create --tenant tenant2 --name "pg1" --pg-
name 10.20.246.23:peerb1 --pg-asn 10.20.246.23,peerb1:4294967295 --pg-bfd-enable
10.20.246.23,peerb1:true --pg-bfd 10.20.246.23,peerb1:30000,30000,50 --pg-next-hop-
self 10.20.246.23,peerb1:true --pg-update-source-ip 10.20.246.23,peerb1:3.3.3.3 --pg-
name 10.20.246.24:peerb1 --pg-asn 10.20.246.24,peerb1:4294967295 --pg-bfd-enable
10.20.246.24,peerb1:true --pg-bfd 10.20.246.24,peerb1:30000,30000,50 --pg-next-hop-
self 10.20.246.24,peerb1:true --pg-update-source-ip 10.20.246.24,peerb1:3.3.3.3

efa tenant service bgp peer create --name B3 --tenant
tenant2 --ipv4-uc-nbr 10.20.246.23,vrf102:121.10.1.3,34566 --ipv4-uc-nbr
```

```
10.20.246.23,vrf102:121.10.2.3,34566 --ipv6-uc-nbr 10.20.246.23,vrf102:121:a::3,34566
--ipv4-uc-nbr 10.20.246.24,vrf102:121.10.1.2,34566 --ipv4-uc-nbr
10.20.246.24,vrf102:121.10.2.2,34566 --ipv6-uc-nbr 10.20.246.24,vrf102:121:a::2,34566


efa tenant epg create --name epg2 --ctag-range  105-107 --po po101 --port
10.20.246.23[0/24] --switchport-mode trunk --tenant tenant2 --vrf vrf101 --anycast-ip
105:11.11.10.254/24 --anycast-ip 106:11.11.1.254/24 --anycast-ip 107:11.11.2.254/24
--anycast-ipv6 105:1001:11:0:1::1/64  --anycast-ipv6 106:1001:11:0:2::1/64 --anycast-
ipv6 107:1002:11:0:3::1/64 --local-ip 105,10.20.246.23:141.10.1.2/24 --local-ip
106,10.20.246.23:141.10.2.2/24 --local-ipv6 107,10.20.246.23:141:a::2/64 --local-ip
105,10.20.246.24:141.10.1.3/24 --local-ip 106,10.20.246.24:141.10.2.3/24 --local-ipv6
107,10.20.246.24:141:a::3/64

efa tenant service bgp peer create --name B2 --tenant
tenant2 --ipv4-uc-nbr 10.20.246.23,vrf101:141.10.1.3,65000 --ipv4-uc-nbr
10.20.246.23,vrf101:141.10.2.3,65000 --ipv6-uc-nbr 10.20.246.23,vrf101:141:a::3,65000
--ipv4-uc-nbr 10.20.246.24,vrf101:141.10.1.2,65000 --ipv4-uc-nbr
10.20.246.24,vrf102:141.10.2.2,65000 --ipv6-uc-nbr 10.20.246.24,vrf102:141:a::2,65000

efa tenant create --name tenant3 --port 10.20.246.23[0/11-14],10.20.246.24[0/11-14] --
vlan-range 30-40 --l2-vni-range 2000-3000  --vrf-count 25 --l3-vni-range 5000-6000

efa tenant po create --name po3 --port 10.20.246.23[0/11],10.20.246.24[0/11] --speed
10Gbps --negotiation active --tenant tenant3

efa tenant vrf create --name vrf31 --tenant tenant3 --rt-type import --rt 301:302 --rt-
type export --rt 305:304
efa tenant vrf create --name vrf32 --tenant tenant3 --rt-type import
--rt 304:305 --rt-type export --rt 300:308 --local-asn 34566 --ipv4-
static-route-next-hop 10.20.246.23,30.0.0.0/24,30.0.0.2 --ipv6-static-route-next-hop
10.20.246.23,5001:1::/64,01::2

efa tenant epg create --name epg3 --ctag-range  30-32 --po po3 --port
10.20.246.23[0/13] --switchport-mode trunk --tenant tenant3 --vrf vrf32 --anycast-
ip 30:30.30.10.254/24 --anycast-ip 32:30.10.1.254/24 --anycast-ip 31:30.10.2.254/24
--anycast-ipv6 30:5001:10:0:1::1/64  --anycast-ipv6 32:5001:10:0:2::1/64 --anycast-
ipv6 31:5002:10:0:3::1/64 --local-ip 30,10.20.246.23:131.10.1.1/24 --local-ip
32,10.20.246.24:131.10.1.2/24 --local-ipv6 31,10.20.246.23:131:a::1/64
efa tenant epg create --name epg32 --ctag-range  35-37 --po po3 --port
10.20.246.23[0/14] --switchport-mode trunk --tenant tenant3 --vrf vrf31 --anycast-
ip 35:11.11.10.254/24 --anycast-ip 36:11.11.1.254/24 --anycast-ip 37:11.11.2.254/24
--anycast-ipv6 35:301:11:0:1::1/64  --anycast-ipv6 36:301:11:0:2::1/64 --anycast-
ipv6 37:302:11:0:3::1/64 --local-ip 35,10.20.246.23:131.10.1.1/24 --local-ip
36,10.20.246.24:131.10.1.2/24 --local-ipv6 37,10.20.246.23:131:a::1/64

efa tenant service bgp peer-group create --tenant tenant3 --name "pg3" --pg-
name 10.20.246.23:peerb3 --pg-asn 10.20.246.23,peerb3:4294967295 --pg-bfd-enable
10.20.246.23,peerb3:true --pg-bfd 10.20.246.23,peerb3:30000,30000,50 --pg-next-hop-
self 10.20.246.23,peerb3:true --pg-update-source-ip 10.20.246.23,peerb3:31.3.3.3 --
pg-name 10.20.246.24:peerb3 --pg-asn 10.20.246.24,peerb3:4294967295 --pg-bfd-enable
10.20.246.24,peerb3:true --pg-bfd 10.20.246.24,peerb3:30000,30000,50 --pg-next-hop-
self 10.20.246.24,peerb3:true --pg-update-source-ip 10.20.246.24,peerb3:3.3.3.3
efa tenant service bgp peer create --name B2 --
tenant tenant3 --ipv4-uc-nbr 10.20.246.23,vrf31:131.10.1.2,4200000000
--ipv4-uc-nbr 10.20.246.24,vrf31:131.10.1.1,4200000000  --ipv6-uc-nbr
10.20.246.24,vrf31:131:a::1,4200000000
```

2. Run the following command to get the operational state of BGP peers belonging to both default VRF and Tenant VRF:

```
efa tenant service bgp peer operational show --tenant <tenant-name> --vrf <vrf-name>
```

3. Run the following command to get the operational state of BGP peers for a given tenant VRF:

```
efa tenant service bgp peer operational show --tenant tenant11 --vrf v1
```

4. Run the following command to get the operational state of BGP peers for a given tenant:

```
efa tenant service bgp peer operational show --tenant tenant11
```

5. Run the following command to get the operational state of BGP peers for all tenant:

```
efa tenant service bgp peer operational show
```

Configure Route Map Attribute

You can configure the route map attribute to enable external connectivity.

About This Task

Follow this procedure to configure the route map attribute when you create or update a BGP peer.

Note

For information about commands and supported parameters to configure route map attribute, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

1. Run the following command to configure route map when you create BGP peer:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
    --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
    --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
    --ipv4-uc-nbr-route-map <device-ip,vrf-name:neighbor-ip,route-mapname,direction(in/
out)>
```

2. Run the following command to configure route map when you update BGP peer:

```
efa tenant service bgp peer update --name <bgp-peer-name> --tenant <tenant-name>
    --operation peer-add
    --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
    --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
```

```
    --ipv4-uc-nbr-route-map <device-ip,vrf-name:neighbor-ip,route-mapname,direction(in/
out)>
```

The following example configures route map attributes:

```
efa tenant service bgp peer update --name ten1bgppeer1 --tenant ten1
    --operation peer-add
    --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
    --ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
    --ipv4-uc-nbr-route-map 10.20.246.15,ten1vrf1:10.20.30.50,rmap2,in
    --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
    --ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.50,true
    --ipv4-uc-nbr-route-map 10.20.246.16,ten1vrf1:10.20.30.50,rmap2,out
```

```
efa tenant service bgp peer show --detail
================================================================================

Name             : ten1bgppeer1
Tenant           : ten1
State            : bgp-peer-created
Description      :

Static Peer
-----------
        Device IP                   : 10.20.246.15
        VRF                         : ten1vrf1
        AFI                         : ipv4
        SAFI                        : unicast
        Remote IP                   : 10.20.30.50
        Remote ASN                  : 50000
        Next Hop Self               : false
        Update Source IP            :
        BFD Enabled                 : true
        BFD Interval                : 0
        BFD Rx                      : 0
        BFD Multiplier              : 0
        MD5 Password                : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS           : false
        Default Originate           : false
        Default Originate Route Map :
        Prefix List In              :
        Prefix List Out             :
        Route Map In                : rmap2
        Route Map Out               : rmap1
        Dev State                   : provisioned
        App State                   : cfg-in-sync

        Device IP                   : 10.20.246.16
        VRF                         : ten1vrf1
        AFI                         : ipv4
        SAFI                        : unicast
        Remote IP                   : fd40:4040:4040:1::fe
        Remote ASN                  : 50000
        Next Hop Self               : false
        Update Source IP            :
        BFD Enabled                 : true
        BFD Interval                : 0
        BFD Rx                      : 0
        BFD Multiplier              : 0
        MD5 Password                : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS           : false
        Default Originate           : false
        Default Originate Route Map :
```

```
        Prefix List In              :
        Prefix List Out             :
        Route Map In                : rmap1
        Route Map Out               : rmap2
        Dev State                   : provisioned
        App State                   : cfg-in-sync


Dynamic Peer
------------
        0 Records
=============================================================================
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.123/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 route-map
in rmap1
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 route-map
in rmap2
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv4 unicast
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.176/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 route-map
out rmap1
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 route-map
out rmap2
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv4 unicast
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Configure remove-private-as on BGP Peer

To enable external connectivity, configure remove-private-as attribute when you create or update BGP peer.

By default, remove-private-as is disabled.

About This Task

Follow this procedure to configure a remove private as.

Procedure

1. Run the following command to create remove private as when you create a BGP Peer on a tenant VRF:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
     --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
     --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
     --ipv4-uc-nbr-remove-private-as <device-ip,vrf-name:neighbor-ip,true|false>
```

2. Run the following command to create remove private as when you update a BGP Peer on a tenant VRF:

```
efa tenant service bgp peer update --name <bgp-peer-name> --tenant <tenant-name>

     --operation peer-add
     --ipv4-uc-nbr 10.20.246.25,v1:10.20.30.50,50000
     --ipv4-uc-nbr-bfd 10.20.246.25,v1:10.20.30.50,true
     --ipv4-uc-nbr-remove-private-as 10.20.246.25,v1:10.20.30.50,true
     --ipv4-uc-nbr 10.20.246.26,v1:10.20.30.50,50000
     --ipv4-uc-nbr-bfd 10.20.246.26,v1:10.20.30.50,true
     --ipv4-uc-nbr-remove-private-as 10.20.246.26,v1:10.20.30.50,false
```

   Example:

```
efa tenant service bgp peer create --name ten1bgppeer1 --tenant tenant11
    --ipv4-uc-nbr 10.20.246.25,v1:10.20.30.40,50000
    --ipv4-uc-nbr-bfd 10.20.246.25,v1:10.20.30.40,true
    --ipv4-uc-nbr-remove-private-as 10.20.246.25,v1:10.20.30.40,true
    --ipv4-uc-nbr 10.20.246.26,v1:10.20.30.40,50000
    --ipv4-uc-nbr-bfd 10.20.246.26,v1:10.20.30.40,true
    --ipv4-uc-nbr-remove-private-as 10.20.246.26,v1:10.20.30.40,true
```

```
10.20.246.25
ORCA_01# show running-config
router bgp
address-family ipv4 unicast vrf v1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 remove-
private-as
  neighbor 10.20.30.40 bfd
  neighbor 10.40.40.253 remote-as
4200000000
  neighbor 10.40.40.253 next-hop-
self
  maximum-paths 8
 !
```

```
10.20.246.26
ORCA_02# show running-config
router bgp
address-family ipv4 unicast vrf v1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 remove-
private-as
  neighbor 10.20.30.40 bfd
  neighbor 10.40.40.252 remote-as
4200000000
  neighbor 10.40.40.252 next-hop-
self
  maximum-paths 8
 !
```

```
efa tenant service bgp peer show --detail
```

```
Name        :
ten1bgppeer1
Tenant      : tenant11
State       : bgp-peer-
created
Description :

Static Peer
-----------
Device
IP             :
10.20.246.25
VRF                  :
 v1
AFI                  :
 ipv4
SAFI                 :
 unicast
Remote
IP             :
10.20.30.40
Remote
ASN            : 50000
Next Hop
Self         : false
Update Source
IP      :
BFD
Enabled          :
true
BFD
Interval         : 0
BFD
Rx               : 0
BFD
Multiplier       : 0
MD5
Password         :
Remove Private
```

```
Device
IP             :
10.20.246.25
VRF                  :
 v1
AFI                  :
 ipv4
SAFI                 :
 unicast
Remote
IP             :
10.20.30.50
Remote
ASN            : 50000
Next Hop
Self         : false
Update Source
IP      :
BFD
Enabled          :
true
BFD
Interval         : 0
BFD
Rx               : 0
BFD
Multiplier       : 0
MD5
Password         :
Remove Private
AS     : true
Default
Originate     : false
Default Originate
Route Map :
Prefix List
In        :
Prefix List
Out       :
```

```
BFD
Multiplier       : 0
MD5
Password         :
Remove Private
AS     : false
Default
Originate     : false
Default Originate
Route Map :
Prefix List
In        :
Prefix List
Out       :
Route Map
In          :
Route Map
Out         :
Dev
State            :
provisioned
App
State            :
cfg-in-sync

Device
IP             :
10.20.246.26
VRF                  :
 v1
AFI                  :
 ipv4
SAFI                 :
 unicast
Remote
IP             :
10.20.30.40
Remote
ASN            : 50000
```

```
AS     : true
Default
Originate     : false
Default Originate
Route Map :
Prefix List
In        :
Prefix List
Out       :
Route Map
In          :
Route Map
Out         :
Dev
State             :
provisioned
App
State             :
cfg-in-sync
```

```
Route Map
In          :
Route Map
Out         :
Dev
State             :
provisioned
App
State             :
cfg-in-sync

Device
IP              :
10.20.246.26
VRF                    :
 v1
AFI                    :
 ipv4
SAFI                   :
 unicast
Remote
IP              :
10.20.30.50
Remote
ASN             : 50000
Next Hop
Self          : false
Update Source
IP      :
BFD
Enabled           :
true
BFD
Interval
 : 0
BFD
Rx
 : 0
```

```
Next Hop
Self          : false
Update Source
IP      :
BFD
Enabled           :
true
BFD
Interval          : 0
BFD
Rx                : 0
BFD
Multiplier        : 0
MD5
Password          :
Remove Private
AS     : true
Default
Originate     : false
Default Originate
Route Map :
Prefix List
In        :
Prefix List
Out       :
Route Map
In          :
Route Map
Out         :
Dev
State             :
provisioned
App
State             :
cfg-in-sync
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.123/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 remove-
private-as
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 remove-
private-as
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.176/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 remove-
private-as
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 remove-
private-as
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Note
For information about commands and supported parameters to configure remove-private-as attribute, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Configure default-originate to Advertise Default Route on BGP Peer

To enable external connectivity, configure the default-originate attribute when you create or update BGP peer (IPv4 and IPv6).

By default, default-originate is disabled.

About This Task

Follow this procedure to configure default originate.

Procedure

1. Run the following command to create a BGP Peer on a tenant VRF:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>

     --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
     --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>

     --ipv4-uc-nbr-default-originate <device-ip,vrf-name:neighbor-ip,true/false>
     --ipv4-uc-nbr-default-originate-route-map <device-ip,vrf-name:neighbor-ip,route-
map>
```

2. Run the following command to update a BGP Peer on a tenant VRF:

```
efa tenant service bgp peer update --name <bgp-peer-name> --tenant <tenant-name>

    --operation peer-add

    --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
    --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>

    --ipv4-uc-nbr-default-originate <device-ip,vrf-name:neighbor-ip,true/false>
    --ipv4-uc-nbr-default-originate-route-map <device-ip,vrf-name:neighbor-ip,route-
map>
```

Note
The `ipv4-uc-nbr-default-originate-route-map` attribute is an optional attribute.

Example:

```
efa tenant service bgp peer create --name ten1bgppeer1 --tenant tenant11
     --ipv4-uc-nbr 10.20.246.3,v1:10.20.30.40,50000
     --ipv4-uc-nbr-bfd 10.20.246.3,v1:10.20.30.40,true
     --ipv4-uc-nbr-default-originate 10.20.246.3,v1:10.20.30.40,true
     --ipv4-uc-nbr 10.20.246.4,v1:10.20.30.40,50000
     --ipv4-uc-nbr-bfd 10.20.246.4,v1:10.20.30.40,true
```

```
    --ipv4-uc-nbr-default-originate 10.20.246.4,v1:10.20.30.40,true
    --ipv4-uc-nbr-default-originate-route-map 10.20.246.4,v1:10.20.30.40,rmap1
```

§efa tenant service bgp peer-group show --detail

```
Name      : ten1bgppeer1
Tenant          :
tenant11
State             : bgp-
peer-created
Description       :

Static Peer
-----------
Device
IP                    :
10.20.246.4
VRF
     : v1
AFI
     : ipv4
SAFI
     : unicast
Remote
IP                    :
10.20.30.40
Remote
ASN                   :
50000
Next Hop
Self                  :
false
Update Source
IP             :
BFD
Enabled
 : true
BFD
Interval
 : 0
BFD
Rx
 : 0
BFD
Multiplier
 : 0
MD5
Password
 :
Remove Private
AS            : false
Default
Originate             :
true
Default Originate
Route Map : rmap1
Prefix List
In               :
Prefix List
Out              :
Route Map
In                 :
Route Map
Device IP           :
10.20.246.4
VRF
     : v1
AFI
     : ipv4
SAFI
     : unicast
Remote
IP                    :
10.20.30.50
Remote
ASN                   :
50000
Next Hop
Self                  :
false
Update Source
IP             :
BFD
Enabled
 : true
BFD
Interval
 : 0
BFD
Rx
 : 0
BFD
Multiplier
 : 0
MD5
Password
 :
Remove Private
AS            : false
Default
Originate             :
false
Default Originate
Route Map :
Prefix List
In               :
Prefix List
Out              :
Route Map
In                 :
Route Map
Out                :
Dev
State
 : provisioned
App
State
 : cfg-in-sync

Device
IP                    :
10.20.246.3
VRF
BFD Multiplier      : 0
MD5
Password
 :
Remove Private
AS            : false
Default
Originate             :
true
Default Originate
Route Map :
Prefix List
In               :
Prefix List
Out              :
Route Map
In                 :
Route Map
Out                :
Dev
State
 : provisioned
App
State
 : cfg-in-sync

Device
IP                    :
10.20.246.3
VRF
     : v1
AFI
     : ipv4
SAFI
     : unicast
Remote
IP                    :
10.20.30.50
Remote
ASN                   :
50000
Next Hop
Self                  :
false
Update Source
IP             :
BFD
Enabled
 : true
BFD
Interval
 : 0
BFD
Rx
 : 0
BFD
Multiplier
 : 0
MD5
Password
```

Out : Dev State : provisioned App State : cfg-in-sync	: v1 AFI : ipv4 SAFI : unicast Remote IP : 10.20.30.40 Remote ASN : 50000 Next Hop Self : false Update Source IP : BFD Enabled : true BFD Interval : 0 BFD Rx : 0	: Remove Private AS : false Default Originate : true Default Originate Route Map : tt Prefix List In : Prefix List Out : Route Map In : Route Map Out : Dev State : provisioned App State : cfg-in-sync

```
efa tenant service bgp peer show --detail
================================================
Name             : bgp173-2501
Tenant           : tenant11
State            : bgp-peer-created
Description      :
Static Peer
-----------
       Device IP                  : 10.20.246.3
    VRF                         : v1
    AFI                         : ipv4
    SAFI                        : unicast
    Remote IP                   : 10.20.30.40
    Remote ASN                  : 50000
    Next Hop Self               : false
    Update Source IP            :
    BFD Enabled                 : true
    BFD Interval                : 0
    BFD Rx                      : 0
    BFD Multiplier              : 0
    MD5 Password                :
    Remove Private AS           : false
    Default Originate           : true
    Default Originate Route Map :
    Prefix List In              :
    Prefix List Out             :
    Route Map In                :
    Route Map Out               :
    Dev State                   : provisioned
    App State                   : cfg-in-sync

      Device IP                  : 10.20.246.4
    VRF                         : v1
    AFI                         : ipv4
    SAFI                        : unicast
```

```
    Remote IP                   : 10.20.30.40
    Remote ASN                  : 50000
    Next Hop Self               : false
    Update Source IP            :
    BFD Enabled                 : true
    BFD Interval                : 0
    BFD Rx                      : 0
    BFD Multiplier              : 0
    MD5 Password                :
    Remove Private AS           : false
    Default Originate           : true
    Default Originate Route Map : rmap1
    Prefix List In              :
    Prefix List Out             :
    Route Map In                :
    Route Map Out               :
    Dev State                   : provisioned
    App State                   : cfg-in-sync
==============================================
```

3. Verify the switch configuration on the SLX device.

10.20.246.3	10.20.246.4

SLX# show running-config router bgp address-family ipv4 unicast vrf v1
  redistribute connected
  neighbor 10.20.30.40 remote-as 50000
  neighbor 10.20.30.40 default-originate
  neighbor 10.20.30.40 bfd
  neighbor 10.40.40.252 remote-as 4200000000
  neighbor 10.40.40.252 next-hop-self
  maximum-paths 8
 !
	SLX# show running-config router bgp address-family ipv4 unicast vrf v1
	  redistribute connected
	  neighbor 10.20.30.40 remote-as 50000
	  neighbor 10.20.30.40 default-originate route-map rmap1
	  neighbor 10.20.30.40 bfd
	  neighbor 10.40.40.253 remote-as 4200000000
	  neighbor 10.40.40.253 next-hop-self
	  maximum-paths 8
	 !

Note
For information about commands and supported parameters to configure default-originate attribute, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Configure Backup Routing Neighbors on BGP Peer

You can configure backup routing neighbors with additional attributes, such as bidirectional forwarding detection, route-maps and prefix-lists. By default, backup routing neighbor is configured with remote-as and next-hop-self. You can provide md5-password as a fabric setting which is applied on the neighbors.

About This Task

Follow this procedure to configure backup routing neighbors on BGP peer.

Note
You can modify the MD5 password for the backup routing neighbors only by configuring and re-configuring the fabric.

Procedure

1. To create backup routing neighbors when you create a BGP peer, run the following command:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
      --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
      --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
      --ipv4-uc-nbr-default-originate <device-ip,vrf-name:neighbor-ip,true/false>
      --ipv4-uc-nbr-default-originate-route-map <device-ip,vrf-name:neighbor-ip,route-
map>
```

2. To create backup routing neighbors when you update a BGP peer, run the following command:

```
efa tenant service bgp peer update --name <bgp-peer-name> --tenant <tenant-name>
      --operation peer-add --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
      --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
      --ipv4-uc-nbr-default-originate <device-ip,vrf-name:neighbor-ip,true/false>
```

3. To remove the backup routing neighbors association with BGP service when you delete a BGP peer, run the following command:

```
efa tenant service bgp peer delete --name <bgp-peer-name> --tenant <tenant-name>
```

Configure Send-Community on Tenant BGP Peer

To enable external connectivity, configure the `send-community` attribute when you create or update a BGP peer (IPv4 and IPv6).

About This Task

Follow this procedure to configure send-community on tenant BGP peer.

Procedure

1. Run the following command to configure send-community when you create a BGP peer:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <tenant-name>
    --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
    --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
    --ipv4-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,
           all|both|extended|large|standard|large-and-standard|large-and-extended>
    --ipv6-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,
           all|both|extended|large|standard|large-and-standard|large-and-extended>
```

2. Run the following command to configure send-community when you update a BGP peer:

```
efa tenant service bgp peer update --name <bgp-peer-name> --tenant <tenant-name>
    --operation peer-add
    --ipv4-uc-nbr <device-ip,vrf-name:neighbor-ip,remote-asn>
    --ipv4-uc-nbr-bfd <switch-ip,vrf-name:ipv4-neighbor,bfd-enable(t/f),bfd-
interval,bfd-rx,bfdmult>
    --ipv4-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,
           all|both|extended|large|standard|large-and-standard|large-and-extended>
```

```
                --ipv6-uc-nbr-send-community <device-ip,vrf-name:neighbor-ip,
            all|both|extended|large|standard|large-and-standard|large-and-extended>
```

Note

`ipv4-uc-nbr-send-community` and `ipv6-uc-nbr-send-community` are an optional attributes.

The following example configures send-community when you create or update a BGP peer:

```
efa tenant service bgp peer create --name ten1bgppeer1 --tenant ten1
    --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
    --ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
    --ipv4-uc-nbr-send-community 10.20.246.15,ten1vrf1:10.20.30.40,all
    --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
    --ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
    --ipv4-uc-nbr-send-community 10.20.246.16,ten1vrf1:10.20.30.40,both

efa tenant service bgp peer update --name ten1bgppeer1 --tenant ten1
    --operation peer-add
    --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
    --ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
    --ipv4-uc-nbr-send-community 10.20.246.15,ten1vrf1:10.20.30.50,standard
    --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
```

```
    --ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.50,true
    --ipv4-uc-nbr-send-community 10.20.246.16,ten1vrf1:10.20.30.50,extended
```

```
efa tenant service bgp peer show --detail
==================================================
Name              : ten1bgppeer1
Tenant            : ten1
State             : bs-state-created
Description       :

Static Peer
-----------
        Device IP         : 10.20.246.15
        VRF               : ten1vrf1
        AFI               : ipv4
        SAFI              : unicast
        Remote IP         : 10.20.30.40
        Remote ASN        : 50000
        Next Hop Self     : false
        Update Source IP  :
        BFD Enabled       : true
        BFD Interval      : 0
        BFD Rx            : 0
        BFD Multiplier    : 0
        Default Originate: 
        Default Originate
        Route Map         :
        Send Community    : all
        Dev State         : provisioned
        App State         : cfg-in-sync

        Device IP         : 10.20.246.15
        VRF               : ten1vrf1
        AFI               : ipv4
        SAFI              : unicast
        Remote IP         : 10.20.30.50
        Remote ASN        : 50000
        Next Hop Self     : false
        Update Source IP  :
        BFD Enabled       : true
        BFD Interval      : 0
        BFD Rx            : 0
        BFD Multiplier    : 0
        Default Originate: false
        Default Originate
        Route Map         :
        Send Community    : standard

        Dev State         : provisioned
        App State         : cfg-in-sync
==================================================
```

```
        Device IP         : 10.20.246.16
        VRF               : ten1vrf1
        AFI               : ipv4
        SAFI              : unicast
        Remote IP         : 10.20.30.40
        Remote ASN        : 50000
        Next Hop Self     : false
        Update Source IP  :
        BFD Enabled       : true
        BFD Interval      : 0
        BFD Rx            : 0
        BFD Multiplier    : 0
        Default Originate: false
        Default Originate
        Route Map         :
        Send Community    : both
        Dev State         : provisioned
        App State         : cfg-in-sync

        Device IP         : 10.20.246.16
        VRF               : ten1vrf1
        AFI               : ipv4
        SAFI              : unicast
        Remote IP         : 10.20.30.50
        Remote ASN        : 50000
        Next Hop Self     : false
        Update Source IP  :
        BFD Enabled       : true
        BFD Interval      : 0
        BFD Rx            : 0
        BFD Multiplier    : 0
        Default Originate: false
        Default Originate
        Route Map         :
        Send Community    : extended
        Dev State         : provisioned
        App State         : cfg-in-sync

Dynamic Peer
-----------
        0 Records
==================================================
```

3. Verify the configuration on SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.123/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 send-
community all
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 send-
community standard
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.176/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  neighbor 10.20.30.40 remote-as
50000
  neighbor 10.20.30.40 send-
community both
  neighbor 10.20.30.40 bfd
  neighbor 10.20.30.50 remote-as
50000
  neighbor 10.20.30.50 send-
community extended
  neighbor 10.20.30.50 bfd
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Configure Out-of-band for a Tenant BGP Peer or Peer Group

You can create out-of-band (OOB) BGP peer group and BGP static or dynamic peer for the use in XCO. Provide the exact BGP peer group or BGP peer configuration in XCO. The configuration enables XCO to manage the BGP peer group and BGP peer created by OOB.

About This Task

Follow this procedure to configure out-of-band BGP peer or peer group for a tenant.

Procedure

1. On both devices, run the **`show running-config router bgp`** command to configure OOB.

2. Run the following command for XCO consumption of OOB BGP Peer Group:

```
(efa:root)root@node-2:~# efa tenant service bgp peer-group create --name ten1bgppg1 --
tenant ten1 --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:65001  --pg-name
10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:65001

Error :  conflicting peer group: [10.20.246.15,pg1:,65001,false,,] and
[10.20.246.15,pg1:,65002,false,,] which is not created by Tenant service

(efa:root)root@node-2:~# efa tenant service bgp peer-group create --name ten1bgppg1 --
tenant ten1 --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:65002  --pg-name
10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:65001

Error :  conflicting peer group: [10.20.246.16,pg1:,65001,false,,] and
[10.20.246.16,pg1:,65002,false,,] which is not created by Tenant service

(efa:root)root@node-2:~# efa tenant service bgp peer-group create --name ten1bgppg1 --
tenant ten1 --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:65002  --pg-name
10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:65002

BgpService created successfully.

(efa:root)root@node-2:~# efa tenant service bgp peer-group show
================================================================================================
=========================
Name        : ten1bgppg1
Tenant      : ten1
State       : bgp-pg-created
+-------------+------+--------+---------+----------+--------+-------------------------
+------------+-----------+
|  Device IP  | Peer | Remote |Next Hop | Update   |  BFD   |          BFD
| Dev State  | App State |
|             | Group|  ASN   |  Self   |Source IP |Enabled |[Interval,Rx,Multiplier]
|            |           |
+-------------+------+--------+---------+----------+--------+-------------------------
+------------+-----------+
| 10.20.246.16|  pg1 | 65002  | false   |          | false  |                         |
provisioned |cfg-in-sync|
+-------------+------+--------+---------+----------+--------+-------------------------
+------------+-----------+
| 10.20.246.15|  pg1 | 65002  | false   |          | false  |                         |
provisioned |cfg-in-sync|
+-------------+------+--------+---------+----------+--------+-------------------------
+------------+-----------+
BGP PeerGroup Details

================================================================================================
=========================
```

3. Run the following command for XCO consumption of OOB BGP Peer.

```
(efa:root)root@node-2:~# efa tenant service bgp peer create --name ten1bgppeer1 --
tenant ten1 --ipv4-uc-dyn-nbr 10.20.246.15,ten1vrf1:10.20.30.0/23,pg1,100 --ipv4-uc-
dyn-nbr 10.20.246.15,ten1vrf1:10.20.40.0/23,pg1,100  --ipv4-uc-dyn-nbr
10.20.246.16,ten1vrf1:10.20.30.0/23,pg1,100 --ipv4-uc-dyn-nbr
10.20.246.16,ten1vrf1:10.20.40.0/23,pg1,50

Error :  conflicting dynamic neighbors: [10.20.246.16,ten1vrf1:10.20.40.0/23,pg1,50]
and [10.20.246.16,ten1vrf1:10.20.40.0/23,pg1,100]

(efa:root)root@node-2:~# efa tenant service bgp peer create --name ten1bgppeer1 --
tenant ten1 --ipv4-uc-dyn-nbr 10.20.246.15,ten1vrf1:10.20.30.0/23,pg1,100 --ipv4-uc-
dyn-nbr 10.20.246.15,ten1vrf1:10.20.40.0/23,pg1,100  --ipv4-uc-dyn-nbr
10.20.246.16,ten1vrf1:10.20.30.0/23,pg1,100 --ipv4-uc-dyn-nbr
10.20.246.16,ten1vrf1:10.20.40.0/23,pg1,100
```

```
BGP Peer created successfully.

(efa:root)root@node-2:~# efa tenant service bgp peer show
================================================================================================
========
Name        : ten1bgppeer1
Tenant      : ten1
State       : bgp-peer-created

+----------+----+----+-----+------+------+--------+---------+-------
+------------------------+-----+-----+
|Device IP |VRF |AFI |SAFI |Remote|Remote|Next Hop| Update  |  BFD  |
BFD           |Dev  |App  |
|          |    |    |     |IP    |ASN   |  Self  |Source IP|Enabled|
[Interval,Rx,Multiplier]|State|State|
+----------+----+----+-----+------+------+--------+---------+-------
+------------------------+-----+-----+
Static Peer Details

+--------------+---------+------+--------+--------------+------+------+------------
+------------+
|  Device-IP   |  VRF    | AFI  |  SAFI  | Listen Range |Listen| Peer | Dev State  |
App State |
|              |         |      |        |              |Limit | Group|
|           |
+--------------+---------+------+--------+--------------+------+------+------------
+------------+
| 10.20.246.16 |ten1vrf1 | ipv4 | unicast| 10.20.30.0/23| 100  | pg1  |provisioned |
cfg-in-sync|
+--------------+---------+------+--------+--------------+------+------+------------
+------------+
| 10.20.246.16 |ten1vrf1 | ipv4 | unicast| 10.20.40.0/23| 100  | pg1  |provisioned |
cfg-in-sync|
+--------------+---------+------+--------+--------------+------+------+------------
+------------+
| 10.20.246.15 |ten1vrf1 | ipv4 | unicast| 10.20.30.0/23| 100  | pg1  |provisioned |
cfg-in-sync|
+--------------+---------+------+--------+--------------+------+------+------------
+------------+
| 10.20.246.15 |ten1vrf1 | ipv4 | unicast| 10.20.40.0/23| 100  | pg1  |provisioned |
cfg-in-sync|
+--------------+---------+------+--------+--------------+------+------+------------
+------------+
Dynamic Peer Details

================================================================================================
========
```

Provision a BGP Peer Group

You can configure a BGP peer group.

About This Task

Complete the following tasks to configure a BGP peer group in your XCO fabric:

Procedure

1. Create a BGP Peer Group on page 327
2. Configure IP Prefix List and Route Map on Tenant BGP Peer Group on page 328
3. Configure Send-Community on Tenant BGP Peer Group on page 332

4. Add Path on Tenant BGP Peer Group on page 335
5. Configure remove-private-as on BGP Peer Group on page 339

Create a BGP Peer Group

You can configure a BGP peer group.

About This Task

Follow this procedure to configure a BGP peer group.

Procedure

1. To create a BGP peer group, run the following command:

```
# efa tenant service bgp peer-group create --name <peer-group-name> --tenant <tenant-
name> --description <description>
    --pg-name <switch-ip:pg-name>
    --pg-asn <switchip:pg-name,remote-asn>
    --pg-bfd <switch-ip:pg-name,bfd-enable(true/false),interval,minrx,multiplier>
    --pg-next-hop-self <switch-ip:pg-name,next-hop-self(true/false/always)>
    -pg-update-source-ip <switch-ip:pg-name,update-source-ip>
```

   The following example creates a BGP peer group:

```
# efa tenant service bgp peer-group create -name ten1BgpPG1 --tenant tenant1
    --pg-name 10.24.80.134:pg1
    --pg-asn 10.24.80.134:pg1,6000
    --pg-bfd 10.24.80.134:pg1,true,100,200,5
    --pg-next-hop-self 10.24.80.134:pg1,true
    --pg-update-source-ip 10.24.80.134:pg1,10.20.30.40
```

2. To update a BGP peer group, run the following command:

```
# efa tenant service bgp peer-group update --name <peer-group-name> --
tenant <tenant-name> --operation <peer-group-add|peer-group-delete|peer-group-desc-
update> --description <description> --pg-name <switch-ip:pg-name> --pg-asn <switch-
ip:pg-name,remote-asn> --pg-bfd <switch-ip:pg-name,bfd-enable(true/false),interval,min-
rx,multiplier> --pg-next-hop-self <switch-ip:pg-name,next-hop-self(true/false/always)>
--pg-update-source-ip <switch-ip:pg-name,update-source-ip>
```

   The following is an example of updating a BGP peer group:

```
efa tenant service bgp peer-group update --name ten1BgpPG1 --tenant tenant1 --
operation peer-group-add --pg-name 10.24.80.134:pg2 -pg-asn 10.24.80.134:pg2,7000 --
pg-bfd 10.24.80.134:pg2,true,200,300,6 --pg-next-hop-self 10.24.80.134:pg2,true --pg-
update-source-ip 10.24.80.134:pg2,10.20.30.41
```

3. To show a BGP peer group, run the following command:

```
efa tenant service bgp peer-group show
================================================================================================
==============
Name : ten1BgpPG1
Tenant : tenant1
State : bs-state-created
+--------------+----------+-------+-----+--------+---------+----+-----------
+------------+-----------+
| Device IP    |PeerGroup | REMOTE| ASN | BFD    | BFD     | BFD| BFD       | Dev-
state  | App-state|
|              |          | ASN   |     | Enabled| Interval| Rx |
Multiplier|           |          |
+--------------+----------+-------+-----+--------+---------+----+-----------
+------------+-----------+
| 10.24.80.134 | pg1      | 6000  | true| 100    | 200     | 5  |provisioned| cfg-in-
```

```
sync|          |
| 10.24.80.134 | pg2      | 7000  | true| 200    | 300     | 6  |provisioned| cfg-in-
sync|          |
+--------------+----------+-------+-----+--------+---------+----+-----------
+------------+----------+
================================================================================================
==============
```

4. To delete a BGP peer group, run the following command:

```
# efa tenant service bgp peer-group delete --name ten1BgpPG1 --tenant tenant1
```

5. Verify the switch configuration on SLX devices.

```
Rack1-Device1# show running-config router bgp
router bgp
local-as 100
neighbor pg1 peer-group
neighbor pg1 remote-as 6000
neighbor pg1 update-source 10.20.30.40
neighbor pg1 next-hop-self neighbor pg1 bfd
neighbor pg1 bfd interval 100 min-rx 200 multiplier 5
neighbor pg2 peer-group
neighbor pg2 remote-as 7000
neighbor pg2 update-source 10.20.30.41
neighbor pg2 next-hop-self neighbor pg2 bfd
neighbor pg2 bfd interval 200 min-rx 300 multiplier 6
address-family ipv4 unicast
!
address-family ipv6 unicast
!
address-family l2vpn evpn
!
```

Configure IP Prefix List and Route Map on Tenant BGP Peer Group

To enable external connectivity, you can configure the IP prefix list and route map attributes in ingress or egress direction when you create or update BGP peer group.

About This Task

Follow this procedure to configure IP prefix list and route map attributes.

Procedure

1. Run the following command to configure IP prefix list and route map attributes when you create BGP peer group:

```
efa tenant service bgp peer-group create --name <bgp-pg-name> --tenant <tenant-name>
    --pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
    --pg-bfd-enable <device-ip,pg-name:true|false>
    --pg-ipv4-uc-nbr-prefix-list <device-ip,pg-name:prefix-list-name,direction>
    --pg-ipv4-uc-nbr-route-map <device-ip,pg-name:route-map-name,direction>

    --pg-ipv6-uc-nbr-prefix-list <device-ip,pg-name:prefix-list-name,direction>
    --pg-ipv6-uc-nbr-route-map <device-ip,pg-name:route-map-name,direction>
```

   The following example configures IP prefix list and route map:

```
efa tenant service bgp peer-group create --name ten1bgppg1 --tenant ten1
    --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:55001
    --pg-bfd-enable 10.20.246.15,pg1:true
    --pg-ipv4-uc-nbr-prefix-list 10.20.246.15,pg1:ipPrefixList1,in
    --pg-ipv6-uc-nbr-prefix-list 10.20.246.15,pg1:ipPrefixList2,out
    --pg-ipv4-uc-nbr-route-map 10.20.246.15,pg1:routeMap2,in
```

```
    --pg-ipv6-uc-nbr-route-map 10.20.246.15,pg1:routeMap1,in

    --pg-name 10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:55001
    --pg-bfd-enable 10.20.246.16,pg1:true
    --pg-ipv4-uc-nbr-prefix-list 10.20.246.16,pg1:ipPrefixList1,in
    --pg-ipv6-uc-nbr-prefix-list 10.20.246.16,pg1:ipPrefixList1,out
    --pg-ipv4-uc-nbr-route-map 10.20.246.16,pg1:routeMap1,in
    --pg-ipv6-uc-nbr-route-map 10.20.246.16,pg1:routeMap1,out
```

2. Run the following command to configure IP prefix list and route map attributes when you update BGP peer group:

```
efa tenant service bgp peer-group update --name <bgp-pg-name> --tenant <tenant-name>
    --operation peer-group-add
    --pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
    --pg-bfd-enable <device-ip,pg-name:true|false>
    --pg-ipv4-uc-nbr-prefix-list <device-ip,pg-name:prefix-list-name,direction>
    --pg-ipv4-uc-nbr-route-map <device-ip,pg-name:route-map-name,direction>

    --pg-ipv6-uc-nbr-prefix-list <device-ip,pg-name:prefix-list-name,direction>
    --pg-ipv6-uc-nbr-route-map <device-ip,pg-name:route-map-name,direction>
```

The following example configures IP prefix list and route map:

```
efa tenant service bgp peer-group update --name ten1bgppg1 --tenant ten1
        --operation peer-group-add
        --pg-name 10.20.246.15:pg2 --pg-asn 10.20.246.15,pg2:55002
        --pg-bfd-enable 10.20.246.15,pg2:true
        --pg-ipv6-uc-nbr-prefix-list 10.20.246.15,pg2:ipPrefixList2,out
        --pg-ipv4-uc-nbr-route-map 10.20.246.15,pg2:routeMap2,in

        --pg-name 10.20.246.16:pg2 --pg-asn 10.20.246.16,pg2:55002
        --pg-bfd-enable 10.20.246.16,pg2:true
```

```
        --pg-ipv6-uc-nbr-prefix-list 10.20.246.16,pg2:ipPrefixList1,out
        --pg-ipv4-uc-nbr-route-map 10.20.246.16,pg2:routeMap1,in
```

```
efa tenant service bgp peer-group
show --detail
=====================================
========
Name              : ten1bgppg1
Tenant            : ten1
State             : bgp-pg-created
Description       :

Peer Group
----------
 Device IP         : 10.20.246.16
 Peer Group        : pg1
 Remote ASN        : 65002
 Next Hop Self     : true
 Update Source IP  : 10.20.30.40
 BFD Enabled       : true
 BFD Interval      : 100
 BFD Rx            : 300
 BFD Multiplier    : 5
 MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS : true
 Prefix List In : Name (afi)
                  ipPrefixList1
(ipv4)
 Prefix List Out   : Name (afi)
                  ipPrefixList1
(ipv6)
 Route Map In      : Name (afi)
                  routeMap1
(ipv4)
 Route Map Out     : Name (afi)
                  routeMap1
(ipv6)
 Send Community    : both (ipv4)
 Dev State         : provisioned
 App State         : cfg-in-sync

 Device IP         : 10.20.246.16
 Peer Group        : pg2
 Remote ASN        : 65002
 Next Hop Self     : true
 Update Source IP  : 10.20.30.50
 BFD Enabled       : true
 BFD Interval      : 100
 BFD Rx            : 300
 BFD Multiplier    : 5
 MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS : true
 Prefix List In    :
 Prefix List Out   : Name (afi)
                  ipPrefixList1
(ipv6)
 Route Map In      : Name (afi)
                  routeMap1
(ipv4)
 Route Map Out     :
 Send Community    : both (ipv4)
```

```
=====================================
=====================




 Device IP         : 10.20.246.15
 Peer Group        : pg1
 Remote ASN        : 65002
 Next Hop Self     : true
 Update Source IP  : 10.20.30.40
 BFD Enabled       : true
 BFD Interval      : 100
 BFD Rx            : 300
 BFD Multiplier    : 5
 MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS : true
 Prefix List In  : Name (afi)
                   ipPrefixList1
(ipv4)
 Prefix List Out   : Name (afi)
                   ipPrefixList1
(ipv6)
 Route Map In      : Name (afi)
                   routeMap1
(ipv6)
                   routeMap2
(ipv4)
 Send Community    : both (ipv4)
 Dev State         : provisioned
 App State         : cfg-in-sync

 Device IP         : 10.20.246.15
 Peer Group        : pg2
 Remote ASN        : 65002
 Next Hop Self     : true
 Update Source IP  : 10.20.30.50
 BFD Enabled       : true
 BFD Interval      : 100
 BFD Rx            : 300
 BFD Multiplier    : 5
 MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS : true
 Prefix List In    :
 Prefix List Out   : Name (afi)
                   ipPrefixList1
(ipv6)
 Route Map In      : Name (afi)
                   routeMap1
(ipv4)
 Route Map Out     :
 Send Community    : both (ipv4)
 Dev State         : provisioned
 App State         : cfg-in-sync
```

```
 Dev State         : provisioned
 App State         : cfg-in-sync
==================================
===
```

```
==================================
============
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor pg1 update-source
10.20.30.40
 neighbor pg1 next-hop-self
 neighbor pg1 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-
rx 300 multiplier 5
 neighbor pg2 peer-group
 neighbor pg2 remote-as 65002
 neighbor pg2 update-source
10.20.30.50
 neighbor pg2 next-hop-self
 neighbor pg2 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor pg2 bfd interval 100 min-
rx 300 multiplier 5
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.214/32
  network 172.31.254.228/32
  neighbor pg2 route-map in
routeMap2
  neighbor pg1 prefix-list
ipPrefixList1 in
  neighbor pg1 route-map in
routeMap2
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
  neighbor pg2 prefix-list
ipPrefixList2 out
  neighbor pg1 prefix-list
ipPrefixList2 out
  neighbor pg1 route-map in
routeMap1
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor pg1 update-source
10.20.30.40
 neighbor pg1 next-hop-self
 neighbor pg1 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-
rx 300 multiplier 5
 neighbor pg2 peer-group
 neighbor pg2 remote-as 65002
 neighbor pg2 update-source
10.20.30.50
 neighbor pg2 next-hop-self
 neighbor pg2 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor pg2 bfd interval 100 min-
rx 300 multiplier 5
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.214/32
  network 172.31.254.246/32
  neighbor pg2 route-map in
routeMap1
  neighbor pg1 prefix-list
ipPrefixList1 in
  neighbor pg1 route-map in
routeMap1
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
  neighbor pg2 prefix-list
ipPrefixList1 out
  neighbor pg1 prefix-list
ipPrefixList1 out
  neighbor pg1 route-map out
routeMap1
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Configure Send-Community on Tenant BGP Peer Group

To enable external connectivity, you can configure the `send-community` attribute when you create or update the BGP peer group.

About This Task

Follow this procedure to configure send-community on tenant BGP peer group.

Procedure

1. Run the following command to configure send-community when you create a BGP peer group:

```
efa tenant service bgp peer-group create --name <bgp-pg-name> --tenant <tenant-name>
    --pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
    --pg-bfd-enable <device-ip,pg-name:true|false>
    --pg-ipv4-uc-nbr-send-community <device-ip,pg-name:
         all|both|extended|large|standard|large-and-extended|large-and-standard>
    --pg-ipv6-uc-nbr-send-community <device-ip,pg-name:
         all|both|extended|large|standard|large-and-extended|large-and-standard>
```

2. Run the following command to configure send-community when you update a BGP peer group:

```
efa tenant service bgp peer-group update --name <bgp-pg-name> --tenant
<tenant-name> --operation peer-group-add --pg-name <device-ip:pg-name> --pg-asn
<device-ip,pg-name:remote-asn> --pg-bfd-enable <device-ip,pg-name:true|false> --pg-
ipv4-uc-nbr-send-community <device-ip,pg-name: all|both|extended|large|standard|large-
and-extended|large-and-standard> --pg-ipv6-uc-nbr-send-community <device-ip,pg-name:
all|both|extended|large|standard|large-and-extended|large-and-standard>
```

The following is an example output of configuring send-community when you create or update a BGP peer group:

```
efa tenant service bgp peer-group create --name ten1bgppg1 --tenant ten1
    --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:55001
    --pg-bfd-enable 10.20.246.15,pg1:true
    --pg-ipv4-uc-nbr-send-community 10.20.246.15,pg1:standard
    --pg-name 10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:55001
    --pg-bfd-enable 10.20.246.16,pg1:true
    --pg-ipv4-uc-nbr-send-community 10.20.246.16,pg1:extended


efa tenant service bgp peer-group update --name ten1bgppg1 --tenant ten1
    --operation peer-group-add
    --pg-name 10.20.246.15:pg2 --pg-asn 10.20.246.15,pg2:55002
    --pg-bfd-enable 10.20.246.15,pg2:true
    --pg-ipv6-uc-nbr-send-community 10.20.246.15,pg2:all
    --pg-name 10.20.246.16:pg2 --pg-asn 10.20.246.16,pg2:55002
```

```
    --pg-bfd-enable 10.20.246.16,pg2:true
    --pg-ipv6-uc-nbr-send-community 10.20.246.16,pg2:both
```

```
efa tenant service bgp peer-group
show --detail
====================================
=========
Name              : ten1bgppg1
Tenant            : ten1
State             : bgp-pg-created
Description       :

Peer Group
----------
  Device IP         : 10.20.246.16
  Peer Group        : pg1
  Remote ASN        : 65002
  Next Hop Self     : true
  Update Source IP  : 10.20.30.40
  BFD Enabled       : true
  BFD Interval      : 100
  BFD Rx            : 300
  BFD Multiplier    : 5
  MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
  Remove Private AS : true
  Prefix List In    :
  Prefix List Out   :
  Route Map In      :
  Route Map Out     :
  Send Community    : standard
(ipv4)
  Dev State         : provisioned
  App State         : cfg-in-sync

  Device IP         : 10.20.246.16
  Peer Group        : pg2
  Remote ASN        : 65002
  Next Hop Self     : true
  Update Source IP  : 10.20.30.50
  BFD Enabled       : true
  BFD Interval      : 100
  BFD Rx            : 300
  BFD Multiplier    : 5
  MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
  Remove Private AS : true
  Prefix List In    :
  Prefix List Out   :
  Route Map In      :
  Route Map Out     :
  Send Community    : all (ipv6)
  Dev State         : provisioned
  App State         : cfg-in-sync
====================================
```

```
====================================
==









  Device IP         : 10.20.246.15
  Peer Group        : pg1
  Remote ASN        : 65002
  Next Hop Self     : true
  Update Source IP  : 10.20.30.40
  BFD Enabled       : true
  BFD Interval      : 100
  BFD Rx            : 300
  BFD Multiplier    : 5
  MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
  Remove Private AS : true
  refix List In     :
  Prefix List Out   :
  Route Map In      :
  Route Map Out     :
  Send Community    : extended
(ipv4)
  Dev State         : provisioned
  App State         : cfg-in-sync

  Device IP         : 10.20.246.15
  Peer Group        : pg2
  Remote ASN        : 65002
  Next Hop Self     : true
  Update Source IP  : 10.20.30.50
  BFD Enabled       : true
  BFD Interval      : 100
  BFD Rx            : 300
  BFD Multiplier    : 5
  MD5 Password      :
$9$QxCvD7N6a0P96eT3BvnQfQ==
  Remove Private AS : true
  Prefix List In    :
  Prefix List Out   :
  Route Map In      :
  Route Map Out     :
  Send Community    : both (ipv6)
  Dev State         : provisioned
  App State         : cfg-in-sync
====================================
```

3. Verify the switch configuration on SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor pg1 update-source
10.20.30.40
 neighbor pg1 next-hop-self
 neighbor pg1 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-
rx 300 multiplier 5
 neighbor pg2 peer-group
 neighbor pg2 remote-as 65002
 neighbor pg2 update-source
10.20.30.50
 neighbor pg2 next-hop-self
 neighbor pg2 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor pg2 bfd interval 100 min-
rx 300 multiplier 5
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.214/32
  network 172.31.254.228/32
  neighbor pg2 prefix-list
ipPrefixList2 in
  neighbor pg2 prefix-list
ipPrefixList2 out
  neighbor pg2 route-map in
routeMap2
  neighbor pg2 route-map out
routeMap2
  neighbor pg1 prefix-list
ipPrefixList1 in
  neighbor pg1 prefix-list
ipPrefixList1 out
  neighbor pg1 route-map in
routeMap1
  neighbor pg1 route-map out
routeMap1
  neighbor pg1 send-community
standard
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
  neighbor pg2 prefix-list
ipPrefixList2 in
  neighbor pg2 prefix-list
ipPrefixList2 out
  neighbor pg2 route-map in
routeMap2
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor pg1 update-source
10.20.30.40
 neighbor pg1 next-hop-self
 neighbor pg1 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-
rx 300 multiplier 5
 neighbor pg2 peer-group
 neighbor pg2 remote-as 65002
 neighbor pg2 update-source
10.20.30.50
 neighbor pg2 next-hop-self
 neighbor pg2 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor pg2 bfd interval 100 min-
rx 300 multiplier 5
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.214/32
  network 172.31.254.246/32
  neighbor pg2 prefix-list
ipPrefixList2 in
  neighbor pg2 prefix-list
ipPrefixList2 out
  neighbor pg2 route-map in
routeMap2
  neighbor pg2 route-map out
routeMap2
  neighbor pg1 prefix-list
ipPrefixList1 in
  neighbor pg1 prefix-list
ipPrefixList1 out
  neighbor pg1 route-map in
routeMap1
  neighbor pg1 route-map out
routeMap1
  neighbor pg1 send-community
extended
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
  neighbor pg2 prefix-list
ipPrefixList2 in
  neighbor pg2 prefix-list
ipPrefixList2 out
  neighbor pg2 route-map in
routeMap2
```

```
  neighbor pg2 route-map out
routeMap2
  neighbor pg2 send-community all
  neighbor pg1 prefix-list
ipPrefixList1 in
  neighbor pg1 prefix-list
ipPrefixList1 out
  neighbor pg1 route-map in
routeMap1
  neighbor pg1 route-map out
routeMap1
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
  neighbor pg2 route-map out
routeMap2
  neighbor pg2 send-community both
  neighbor pg1 prefix-list
ipPrefixList1 in
  neighbor pg1 prefix-list
ipPrefixList1 out
  neighbor pg1 route-map in
routeMap1
  neighbor pg1 route-map out
routeMap1
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Add Path on Tenant BGP Peer Group

You can add paths on tenant BGP peer group.

About This Task

Follow this procedure to configure additional paths (for both IPv4 and IPv6) when you create or update a BGP peer-group.

Procedure

1. To configure an additional path when you create a BGP peer-group, run the following command:

```
efa tenant service bgp peer-group create --name <bgp-pg-name> --tenant <tenant-name>
   --pg-ipv4-uc-nbr-add-path-capability <device-ip,pg-name:{send | receive | both}>
   --pg-ipv4-uc-nbr-add-path-advertise-all <device-ip,pg-name:{true | false}>
   --pg-ipv4-uc-nbr-add-path-advertise-group-best <device-ip,pg-name:{true | false}>
   --pg-ipv4-uc-nbr-add-path-advertise-best <device-ip,pg-name: 2-16>
```

2. To configure an additional path when you update a BGP peer-group, run the following command:

```
efa tenant service bgp peer-group update --name <bgp-pg-name> --tenant <tenant-name>
   --pg-ipv4-uc-nbr-add-path-capability <device-ip,pg-name:{send | receive | both}>
   -pg-ipv4-uc-nbr-add-path-advertise-all <device-ip,pg-name:{true | false}>
   -pg-ipv4-uc-nbr-add-path-advertise-group-best <device-ip,pg-name:{true | false}>
   -pg-ipv4-uc-nbr-add-path-advertise-best <device-ip,pg-name: 2-16>
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor pg1 update-source
10.20.30.40
 neighbor pg1 next-hop-self
 neighbor pg1 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-
rx 300 multiplier 5
 neighbor pg2 peer-group
 neighbor pg2 remote-as 65002
 neighbor pg2 update-source
10.20.30.50
 neighbor pg2 next-hop-self
 neighbor pg2 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor pg2 bfd interval 100 min-
rx 300 multiplier 5
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.214/32
  network 172.31.254.228/32
  additional-paths select all
  neighbor pg1 additional-paths
send receive
  neighbor pg1 additional-paths
advertise best 10 group-best all
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
  additional-paths select all
  neighbor pg1 additional-paths
send
  neighbor pg1 additional-paths
advertise best 8 group-best all
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 65002
 neighbor pg1 update-source
10.20.30.40
 neighbor pg1 next-hop-self
 neighbor pg1 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-
rx 300 multiplier 5
 neighbor pg2 peer-group
 neighbor pg2 remote-as 65002
 neighbor pg2 update-source
10.20.30.50
 neighbor pg2 next-hop-self
 neighbor pg2 password
$9$QxCvD7N6a0P96eT3BvnQfQ==
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor pg2 bfd interval 100 min-
rx 300 multiplier 5
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.214/32
  network 172.31.254.246/32
  additional-paths select all
  neighbor pg1 additional-paths
send receive
  neighbor pg1 additional-paths
advertise best 5
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
  additional-paths select all
  neighbor pg1 additional-paths
receive
  neighbor pg1 additional-paths
advertise best 4
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Example

The following is an example output for adding an additional paths when you create or update a BGP peer group:

```
efa tenant service bgp peer-group create --name ten1bgppg1 --tenant ten1

       --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:55001
       --pg-bfd-enable 10.20.246.15,pg1:true
       --pg-ipv4-uc-nbr-add-path-capability 10.20.246.15,pg1:both
       --pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,pg1:true
       --pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,pg1:true
       --pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,pg1:10

       --pg-name 10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:55001
       --pg-bfd-enable 10.20.246.16,pg1:true
       --pg-ipv4-uc-nbr-add-path-capability 10.20.246.16,pg1:both
       --pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.16,pg1:false
       --pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.16,pg1:false
       --pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.16,pg1:5


efa tenant service bgp peer-group update --name ten1bgppg1 --tenant ten1
       --operation peer-group-add
       --pg-name 10.20.246.15:pg2 --pg-asn 10.20.246.15,pg2:55002
       --pg-bfd-enable 10.20.246.15,pg2:true
       --pg-ipv4-uc-nbr-add-path-capability 10.20.246.15,pg2:send
       --pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,pg2:true
       --pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,pg2:true
       --pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,pg2:8

       --pg-name 10.20.246.16:pg2 --pg-asn 10.20.246.16,pg2:55002
       --pg-bfd-enable 10.20.246.16,pg2:true
       --pg-ipv4-uc-nbr-add-path-capability 10.20.246.16,pg2:receive
       --pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.16,pg2:false
```

```
      --pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.16,pg2:false
      --pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.16,pg2:4
```

```
efa tenant service bgp peer-group
show --detail
==================================
======================
Name              : ten1bgppg1
Tenant            : ten1
State             : bgp-pg-created
Description       :

Peer Group
----------
 Device IP                :
10.20.246.16
 Peer Group               : pg1
 Remote ASN               : 65002
 Next Hop Self            : true
 Update Source IP         :
10.20.30.40
 BFD Enabled              : true
 BFD Interval             : 100
 BFD Rx                   : 300
 BFD Multiplier           : 5
 MD5 Password             :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS        : true
 Add Path Capability IPv4 : Send,
Receive
 Add Path Advertise       : Best 5
 Dev State                :
provisioned
 App State                : cfg-in-
sync

 Device IP                :
10.20.246.16
 Peer Group               : pg2
 Remote ASN               : 65002
 Next Hop Self            : true
 Update Source IP         :
10.20.30.50
 BFD Enabled              : true
 BFD Interval             : 100
 BFD Rx                   : 300
 BFD Multiplier           : 5
 MD5 Password             :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS        : true
 Add Path Capability      : Receive
 Add Path Advertise       : Best 4
 Dev State                :
provisioned
 App State                : cfg-in-
sync
==================================
=======
```

```
==================================
======================




 Device IP                :
10.20.246.15
 Peer Group               : pg1
 Remote ASN               : 65002
 Next Hop Self            : true
 Update Source IP         :
10.20.30.40
 BFD Enabled              : true
 BFD Interval             : 100
 BFD Rx                   : 300
 BFD Multiplier           : 5
 MD5 Password             :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS        : true
 Add Path Capability IPv4 : Send,
Receive
 Add Path Advertise  IPv4 : All,
Group Best, Best 10
 Dev State                :
provisioned
 App State                : cfg-in-
sync

 Device IP                :
10.20.246.15
 Peer Group               : pg2
 Remote ASN               : 65002
 Next Hop Self            : true
 Update Source IP         :
10.20.30.50
 BFD Enabled              : true
 BFD Interval             : 100
 BFD Rx                   : 300
 BFD Multiplier           : 5
 MD5 Password             :
$9$QxCvD7N6a0P96eT3BvnQfQ==
 Remove Private AS        : true
 Add Path Capability IPv6 : Send
 Add Path Advertise  IPv6 : All,
Group Best, Best 8
 Dev State                :
provisioned
 App State                : cfg-in-
sync
==================================
=====
```

Configure remove-private-as on BGP Peer Group

To enable external connectivity, configure the remove-private-as attribute when you create or update BGP peer group.

By default, remove-private-as is disabled.

About This Task

Follow this procedure to configure remove-private-as.

Procedure

1. Run the following command to configure a remove-private-as when you create a BGP Peer-Group on a tenant VRF:

```
efa tenant service bgp peer-group create --name <bgp-pg-name> --tenant <tenant-name>
    --pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
    --pg-bfd-enable <device-ip,pg-name:true|false>
    --pg-remove-private-as <device-ip,pg-name:true|false>
```

2. Run the following command to configure a remove-private-as when you update a BGP Peer-Group on a tenant VRF:

```
efa tenant service bgp peer-group update --name <bgp-pg-name> --tenant <tenant-name>
    --operation peer-group-add
    --pg-name <device-ip:pg-name> --pg-asn <device-ip,pg-name:remote-asn>
    --pg-bfd-enable <device-ip,pg-name:true|false>
    --pg-remove-private-as <device-ip,pg-name:true|false>
```

Example:

```
efa tenant service bgp peer-group create --name ten1bgppg1 --tenant ten1
    --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:55001
    --pg-bfd-enable 10.20.246.15,pg1:true
    --pg-remove-private-as 10.20.246.15,pg1:true
    --pg-name 10.20.246.16:pg1 --pg-asn 10.20.246.16,pg1:55001
    --pg-bfd-enable 10.20.246.16,pg1:true
    --pg-remove-private-as 10.20.246.16,pg1:true

efa tenant service bgp peer-group update --name ten1bgppg1 --tenant ten1
    --operation peer-group-add
    --pg-name 10.20.246.15:pg2 --pg-asn 10.20.246.15,pg2:55002
    --pg-bfd-enable 10.20.246.15,pg2:true
    --pg-remove-private-as 10.20.246.15,pg2:true
    --pg-name 10.20.246.16:pg2 --pg-asn 10.20.246.16,pg2:55002
    --pg-bfd-enable 10.20.246.16,pg2:true
    --pg-remove-private-as 10.20.246.16,pg2:true

efa tenant service bgp peer-group show --detail
================================================
Name             : ten1bgppg1
Tenant           : ten1
State            : bgp-pg-state-created

Peer Group
----------
        Device IP        : 10.20.246.15
        Peer Group       : pg1
        Remote ASN       : 55001
        Next Hop Self    : false
        BFD Enabled      : true
        BFD Interval     :
        BFD Rx           :
        BFD Multiplier   :
```

```
        Remove Private AS: true
        Dev State        : provisioned
        App State        : cfg-in-sync

        Device IP        : 10.20.246.15
        Peer Group       : pg2
        Remote ASN       : 55002
        Next Hop Self    : false
        BFD Enabled      : true
        BFD Interval     :
        BFD Rx           :
        BFD Multiplier   :
        Remove Private AS: true
        Dev State        : provisioned
        App State        : cfg-in-sync

        Device IP        : 10.20.246.16
        Peer Group       : pg1
        Remote ASN       : 55001
        Next Hop Self    : false
        BFD Enabled      : true
        BFD Interval     :
        BFD Rx           :
        BFD Multiplier   :
        Remove Private AS: true
        Dev State        : provisioned
        App State        : cfg-in-sync

        Device IP        : 10.20.246.16
        Peer Group       : pg2
        Remote ASN       : 55002
        Next Hop Self    : false
        BFD Enabled      : true
        BFD Interval     :
        BFD Rx           :
        BFD Multiplier   :
        Remove Private AS: false
        Dev State        : provisioned
        App State        : cfg-in-sync
================================================
```

3. Verify the switch configuration on the SLX device.

```
Rack1-Device1# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 55001
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg2 peer-group
 neighbor pg2 remote-as 55002
 neighbor pg2 remove-private-as
 neighbor pg2 bfd
 neighbor 10.20.20.4 remote-as
4200000000
 neighbor 10.20.20.4 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.123/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

```
Rack1-Device2# show running-config
router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 55001
 neighbor pg1 remove-private-as
 neighbor pg1 bfd
 neighbor pg2 peer-group
 neighbor pg2 remote-as 55002
 neighbor pg2 password remove-
private-as
 neighbor pg2 bfd
 neighbor 10.20.20.5 remote-as
4200000000
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.46/32
  network 172.31.254.176/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf
ten1vrf1
  redistribute connected
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

Note
For information about commands and supported parameters to configure remove-private-as attribute, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Activate Peer Group on Tenant BGP

Activate BGP peer group under the IPv6 unicast address family of the default VRF.

About This Task

Follow this procedure to activate a tenant BGP peer group.

XCO automatically activates the BGP peer group under the IPv6 unicast address family of the default VRF when you create a BGP peer group. The BGP peer group remains activated till the BGP peer group remains instantiated.

Note

An upgrade to XCO 3.2.1 or later automatically activates all existing peer groups, and the corresponding BGP sessions are reset for the “activate” configuration to be effective.

Procedure

Run the following command to activate a BGP peer group:

```
efa tenant service bgp peer-group <create/update>
```

Example

```
efa tenant service bgp peer-group create --name ten1BgpPG1--tenant tenant1
          --pg-name 10.20.246.15:pg7 --pg-asn 10.20.246.15,pg7:6001
          --pg-update-source-ip 10.20.246.15,pg7:10.20.30.40 --pg-bfd-enable
10.20.246.15,pg7:true
          --pg-bfd 10.20.246.15,pg7:100,200,5 --pg-next-hop-self 10.20.246.15,pg7:true

efa tenant service bgp peer-group update --name ten1BgpPG1 --tenant tenant1 --operation
peer-group-add
          --pg-name 10.20.246.15:pg8 --pg-asn 10.20.246.15,pg8:7001
          --pg-update-source-ip 10.20.246.15,pg8:10.20.30.41 --pg-bfd-enable
10.20.246.15,pg8:true
          --pg-bfd 10.20.246.15,pg8:200,300,6 --pg-next-hop-self 10.20.246.15,pg8:true


efa tenant service bgp peer-group show --tenant tenant1
================================================================================================
=================================================================
Name        : ten1BgpPG1
Tenant      : tenant1
State       : bgp-pg-created

+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
|  Device IP   | Peer  | Remote | Next Hop |   Update    |   BFD   |
BFD            |  Dev State  |  App State  |
|              | Group |  ASN   |   Self   |  Source IP  | Enabled |
[Interval,Rx,Multiplier] |             |             |
+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
| 10.20.246.15 | pg7   |  6001  |   true   | 10.20.30.40 |  true   |       100, 200,
5        | provisioned | cfg-in-sync |
+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
| 10.20.246.15 | pg8   |  7001  |   true   | 10.20.30.41 |  true   |       200, 300,
6        | provisioned | cfg-in-sync |
+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
BGP PeerGroup Details


Rack1-Device1# show running-config router bgp
router bgp
local-as 4200000000
capability as4-enable
fast-external-fallover
```

```
neighbor pg7 peer-group
neighbor pg7 remote-as 6001
neighbor pg7 update-source 10.20.30.40
neighbor pg7 next-hop-self
neighbor pg7 bfd
neighbor pg7 bfd interval 100 min-rx 200 multiplier 5
neighbor pg8 peer-group
neighbor pg8 remote-as 7001
neighbor pg8 update-source 10.20.30.41
neighbor pg8 next-hop-self
neighbor pg8 bfd
neighbor pg8 bfd interval 200 min-rx 300 multiplier 6
neighbor 10.20.20.3 remote-as 4200000000
neighbor 10.20.20.3 next-hop-self
address-family ipv4 unicast
  network 172.31.254.55/32
  network 172.31.254.152/32
  maximum-paths 8
  graceful-restart
!
address-family ipv6 unicast
  neighbor pg7 activate
  neighbor pg8 activate
!
address-family l2vpn evpn
  graceful-restart
!
!
Rack1-Device1#



efa tenant service bgp peer-group update --name ten1BgpPG1 --tenant tenant1
                      --operation peer-group-delete --pg-name 10.20.246.15:pg8

efa tenant service bgp peer-group show --tenant tenant1
================================================================================================
=====================================================================
Name        : ten1BgpPG1
Tenant      : tenant1
State       : bgp-pg-created

+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
|  Device IP   | Peer  | Remote | Next Hop |   Update    |   BFD   |
BFD            |  Dev State  |  App State  |
|              | Group |  ASN   |   Self   |  Source IP  | Enabled |
[Interval,Rx,Multiplier] |            |             |
+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
| 10.20.246.15 | pg7   |  6001  |   true   | 10.20.30.40 |  true   |       100, 200,
5        | provisioned | cfg-in-sync |
+--------------+-------+--------+----------+-------------+---------
+--------------------------+-------------+-------------+
BGP PeerGroup Details


Rack1-Device1# show running-config router bgp
router bgp
local-as 4200000000
capability as4-enable
fast-external-fallover
neighbor pg7 peer-group
neighbor pg7 remote-as 6001
```

```
neighbor pg7 update-source 10.20.30.40
neighbor pg7 next-hop-self
neighbor pg7 bfd
neighbor pg7 bfd interval 100 min-rx 200 multiplier 5
neighbor 10.20.20.3 remote-as 4200000000
neighbor 10.20.20.3 next-hop-self
address-family ipv4 unicast
  network 172.31.254.55/32
  network 172.31.254.152/32
  maximum-paths 8
  graceful-restart
!
address-family ipv6 unicast
  neighbor pg7 activate
!
address-family l2vpn evpn
  graceful-restart
!
!
Rack1-Device1#


efa tenant service bgp peer-group delete --name ten1BgpPG1 --tenant tenant1


Rack1-Device1# show running-config router bgp
router bgp
local-as 4200000000
capability as4-enable
fast-external-fallover
neighbor 10.20.20.3 remote-as 4200000000
neighbor 10.20.20.3 next-hop-self
address-family ipv4 unicast
  network 172.31.254.55/32
  network 172.31.254.152/32
  maximum-paths 8
  graceful-restart
!
address-family ipv6 unicast
!
address-family l2vpn evpn
  graceful-restart
!
!
Rack1-Device1#
```

Share Resources Across Tenants using Shared Tenant

One tenant, with the `role=shared` attribute, owns the resources and entities that can be shared across all the other tenants, called non-Shared Tenant. The tenant service can have one Shared Tenant that services all the shared resources. The Shared Tenant owns the physical ports, Layer 2 and Layer 3 VNI number ranges, VLAN number ranges, and the VRF numbers. The Shared Tenant can create the endpoints and the VRFs, but not the endpoint groups.

A non-Shared Tenant cannot use the ports that the Shared Tenant owns if the ports are already part of an endpoint. A non-Shared Tenant cannot create an endpoint using the ports that the Shared Tenant owns.

Example: Shared Port use case (Layer 2 hand-off)

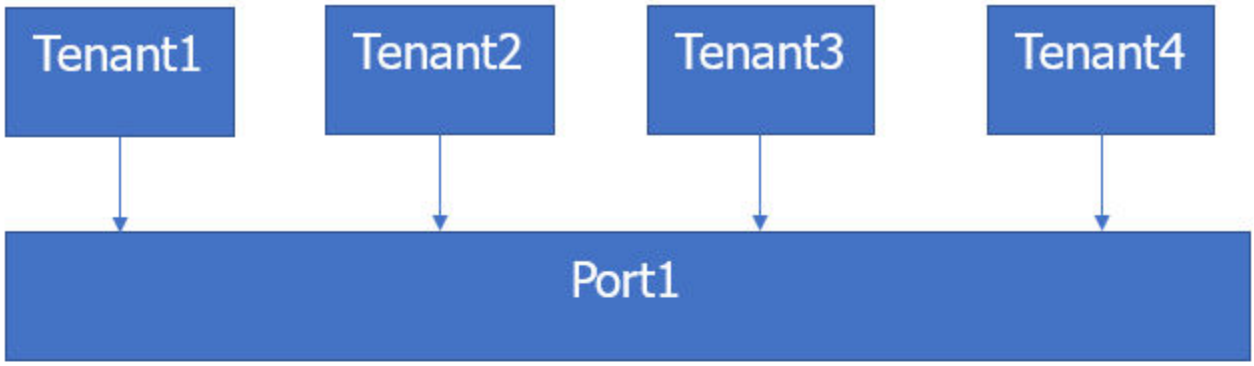

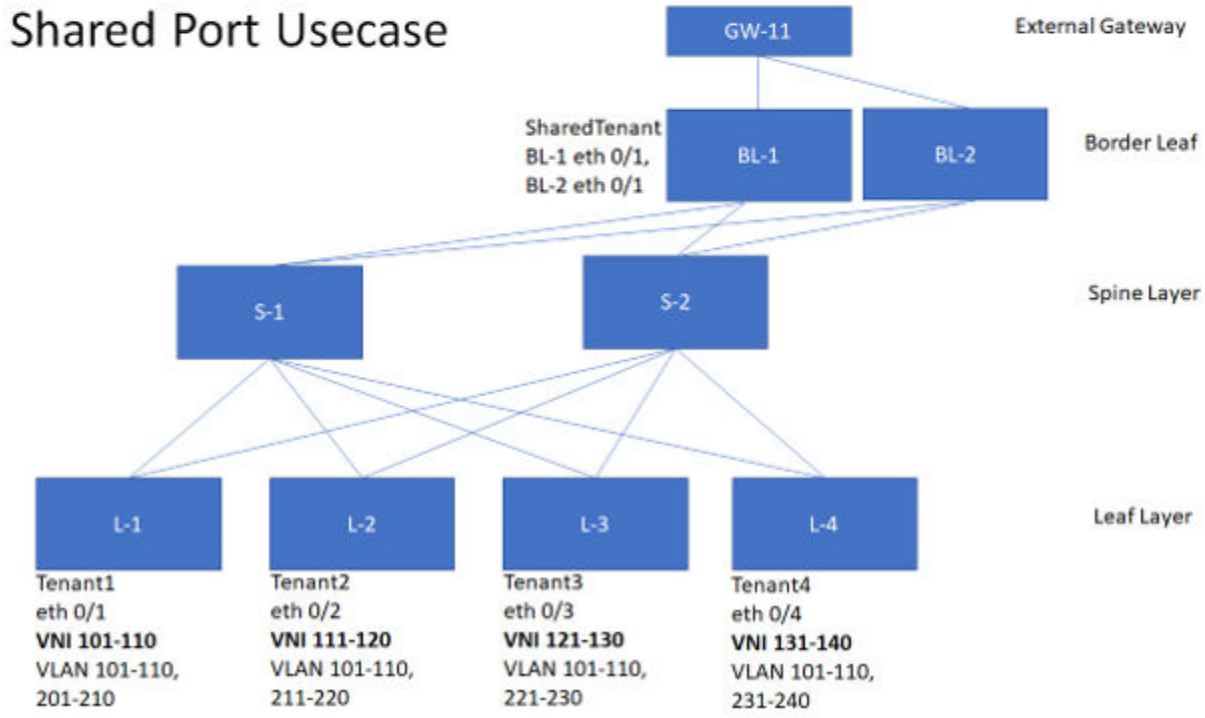

The following examples show the commands and syntax used to configure the Shared Port.

```
efa tenant create --name tenant1 --l2-vni-range 101-110 --vlan-range 101-110,201-210
--port L-1[0/1]
efa tenant create --name tenant2 --l2-vni-range 111-120 --vlan-range 101-110,211-220
--port L-2[0/2]
efa tenant create --name tenant3 --l2-vni-range 121-130 --vlan-range 101-110,221-230
--port L-3[0/3]
efa tenant create --name tenant4 --l2-vni-range 131-140 --vlan-range 101-110,231-240
--port L-4[0/4]
```

```
efa tenant create --name SharedTenant --port BL-1[0/1],BL-2[0/1] --type shared
```

```
efa tenant epg create --name ten1epg1 --tenant tenant1 --port L-1[0/1]
--switchport-mode trunk --ctag-range 101-110 --l2-vni 101:101 --l2-vni 102:102
--l2-vni 110:110

efa tenant epg create --name ten2epg1 --tenant tenant2 --port L-2[0/2]
--switchport-mode trunk --ctag-range 101-110 --l2-vni 101:111 --l2-vni 102:112
--l2-vni 110:120

efa tenant epg create --name ten3epg1 --tenant tenant3 --port L-3[0/3]
--switchport-mode trunk --ctag-range 101-110 --l2-vni 101:121 --l2-vni 102:122
--l2-vni 110:130

efa tenant epg create --name ten4epg1 --tenant tenant4 --port L-4[0/4]
--switchport-mode trunk --ctag-range 101-110 --l2-vni 101:131 --l2-vni 102:132
--l2-vni 110:140
```

```
efa tenant epg create --name ten1epg2 --tenant tenant1 --port BL-1[0/1],BL-2[0/1]
--switchport-mode trunk --ctag-range 201-210 --l2-vni 201:101 --l2-vni 202:102
--l2-vni 210:110

efa tenant epg create --name ten2epg2 --tenant tenant2 --port BL-1[0/1],BL-2[0/1]
--switchport-mode trunk --ctag-range 211-220 --l2-vni 211:111 --l2-vni 212:112
--l2-vni 220:120

efa tenant epg create --name ten3epg2 --tenant tenant3 --port BL-1[0/1],BL-2[0/1]
--switchport-mode trunk --ctag-range 221-230 --l2-vni 221:121 --l2-vni 212:122
--l2-vni 230:130

efa tenant epg create --name ten4epg2 --tenant tenant4 --port BL-1[0/1],BL-2[0/1]
--switchport-mode trunk --ctag-range 231-240 --l2-vni 231:131 --l2-vni 212:132
--l2-vni 240:140
```

Example: Shared Endpoint use case (Layer 2 hand-off)

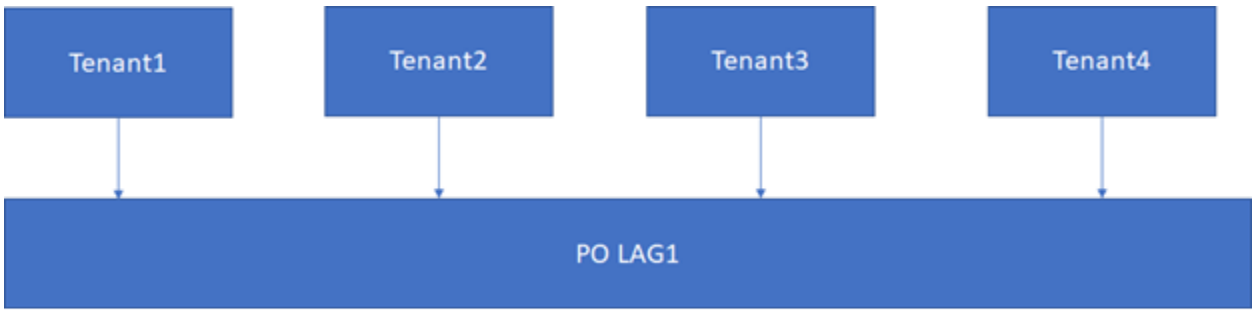

The following examples show the commands and syntax used to configure the Shared Endpoint.

```
efa tenant create --name SharedTenant --port BL-1[0/1],BL-2[0/1]
--l3-vni-range 1001-1010 --vrf-count 10 --type shared efa tenant vrf create --name red
--tenant SharedTenant

efa tenant epg create --name ten1epg1 --tenant tenant1 --port L-1[0/1]
--switchport-mode trunk --ctag-range 101-102 --l2-vni 101:101 --l2-vni 102:102
--anycast-ip 101:10.10.10.1/24 --vrf red --l3-vni 1001

efa tenant epg create --name ten2epg1 --tenant tenant2 --port L-2[0/2]
--switchport-mode trunk --ctag-range 101-102 --l2-vni 101:111 --l2-vni 102:112
--anycast-ip 101:10.10.11.1/24 --vrf red --l3-vni 1001

efa tenant epg create --name ten3epg1 --tenant tenant3 --port L-3[0/3]
--switchport-mode trunk --ctag-range 101-102 --l2-vni 101:121 --l2-vni 102:122
--anycast-ip 101:10.10.12.1/24 --vrf red --l3-vni 1001

efa tenant epg create --name ten1epg1 --tenant tenant4 --port L-4[0/4]
--switchport-mode trunk --ctag-range 101-102 --l2-vni 101:131 --l2-vni 102:132
--anycast-ip 101:10.10.13.1/24 --vrf red --l3-vni 1001
```

Example: Shared Endpoint use case (Layer 3 hand-off)

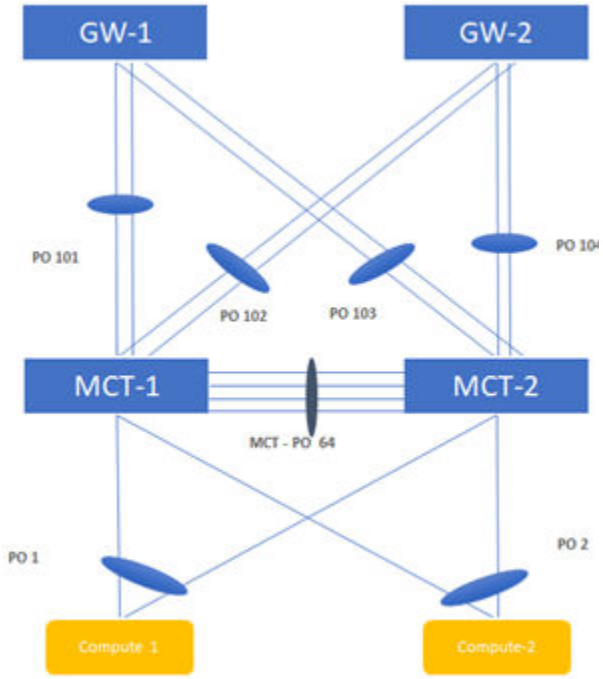

Figure 23: Topology

The following examples show the commands and syntax used to configure the Shared Endpoint.

```
efa tenant create --name tenant1 --l2-vni-range 1001-1010 --vlan-range 1001-1010
--port BL-1[0/11],BL-2[0/11] --l3-vni-range 10001-10010 --vrf-count 10

efa tenant create --name tenant2 --l2-vni-range 1101-1110 --vlan-range 1101-1110
--port BL-1[0/21],BL-2[0/21] --l3-vni-range 20001-20010 --vrf-count 10
```

```
efa tenant vrf create --name vrf1 --tenant Tenant1
efa tenant vrf create --name vrf2 --tenant Tenant2
```

```
efa tenant epg create --name ten1epg1 --tenant tenant1 --port BL-1[0/11]
--switchport-mode trunk --ctag-range 1001 --l2-vni 1001:1001 --anycast-ip
1001:10.10.10.1/24
--vrf vrf1 --l3-vni 1001
efa tenant epg create --name ten2epg1 --tenant tenant2 --port BL-1[0/21]
--switchport-mode trunk --ctag-range 1101 --l2-vni 1101:1101 --anycast-ip
1101:10.10.11.1/24
--vrf vrf2 --l3-vni 1002
```

```
efa tenant create --name SharedTenant --port BL-1[0/1-8],BL-2[0/1-8] --type shared
```

```
efa tenant po create --name po101 --tenant SharedTenant --speed 10Gbps
--negotiation active --port BL-1[0/1],BL-1[0/2]
efa tenant po create --name po102 --tenant SharedTenant --speed 10Gbps
--negotiation active --port BL-1[0/3],BL-1[0/4]
```

VRF1

```
efa tenant epg create --name ten1epg2 --tenant tenant1 --type l3-handover
--po po101 --switchport-mode trunk --ctag-range 101 --vrf vrf1  --local-ipv4-address
11.1.1.1/30
--local-ipv6-address 2001:11:1:1::1/126 --remote-ipv4-address 11.1.1.2 --remote-ipv6-
address
2001:11:1:1::2 --remote-as 4220000001 --bfd --bfd-interval 100 --bfd-min-rx 200 --bfd-
multiplier 10
```

```
efa tenant epg create --name ten1epg3 --tenant tenant1 --type l3-handover
--po po102 --switchport-mode trunk --ctag-range 201 --vrf vrf1 --local-ipv4-address
12.1.1.1/30
--local-ipv6-address 2001:12:1:1::1/126 --remote-ipv4-address 12.1.1.2 --remote-ipv6-
address
2001:12:1:1::2 --remote-as 4220000001 --bfd --bfd-interval 100 --bfd-min-rx 200 --bfd-
multiplier 10
```

VRF2

```
efa tenant epg create --name ten2epg2 --tenant tenant2 --type l3-handover --po po101
--switchport-mode trunk --ctag-range 102 --vrf vrf2  --local-ipv4-address 11.2.1.1/30
--local-ipv6-address 2001:11:2:1::1/126 --remote-ipv4-address 11.2.1.2 --remote-ipv6-
address
2001:11:1:1::2 --remote-as 4220000001 --bfd --bfd-interval 100 --bfd-min-rx 200 --bfd-
multiplier 10
```

```
efa tenant epg create --name ten2epg3 --tenant tenant2 --type l3-handover --po po102
--switchport-mode trunk --ctag-range 202 --vrf vrf2 --local-ipv4-address 12.2.1.1/30
--local-ipv6-address 2001:12:2:1::1/126 --remote-ipv4-address 12.2.1.2 --remote-ipv6-
address
2001:12:2:1::2 --remote-as 4220000001 --bfd  --bfd-interval 100 --bfd-min-rx 200 --bfd-
multiplier 10
```

Shared VRF and Router

XCO provides a provisioning model to support sharing of VRF or Router across multiple tenants. The following models are supported:

- Inter-POD (Inter-Tenant) routing (Tenant1 routing to Tenant2 and vice versa)
- Multiple tenants accessing a shared resource (for example, storage)
- Multiple tenants using a shared VRF for l3-hand-off

Entities (VRFs) created by the shared tenant are available for the use by all the Private Tenants.

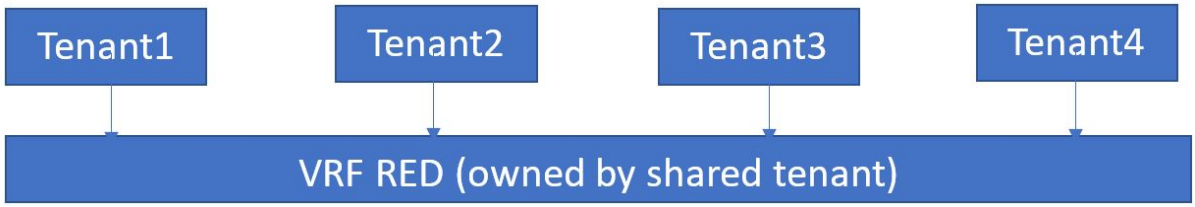

Configure Shared Tenant, Shared VRF, and Private EPG using Shared VRF

You can configure shared tenant, shared VRF, and private EPG.

About This Task

Follow this procedure to configure shared tenant, shared VRF, and private EPG.

Procedure

1. To configure Shared Tenant, run the following command:

```
efa tenant create --name <epg-name> --type shared --port <port-list>
    --vrf-count <num-of-vrfs> --l3-vni-range <l3-vni-range>
    --vlan-range <vlan-range> --l2-vni-range <l2-vni-range>
```

2. To configure Shared VRF under the ownership of Shared Tenant, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <shared-tenant-name>
```

3. To configure endpoint group (EPG) under the ownership of Private Tenant using the Shared VRF, run the following command:

```
efa tenant epg create --name <epg-name> --tenant <private-tenant>
    --po <po-list> --switchport-mode <trunk|access> --ctag-range <ctag-range>
    --anycast-ip <ctag:anycast-ip> --vrf <shared-vrf-owned-by-shared-tenant>
```

Configure L3-Hand-Off EPG and BGP Peer under Ownership of Shared Tenant

L3-Hand-Off endpoint groups (EPG) created under the ownership of Shared Tenant are exclusively meant for the hand-off of the Shared VRF. You cannot create the Extension EPGs under the ownership of Shared Tenant. BGP peer created under the ownership of Shared Tenant are exclusively meant for the hand-off of the Shared VRF.

Before You Begin

Follow this procedure to configure L3-hand-off EPG and BGP peer.

Procedure

1. To configure shared tenant, run the following command:

```
efa tenant create --name <epg-name> --type shared --port <port-list>
    --vrf-count <num-of-vrfs> --l3-vni-range <l3-vni-range>
    --vlan-range <vlan-range> --l2-vni-range <l2-vni-range>
```

2. To configure shared VRF under the ownership of shared tenant, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <shared-tenant-name>
```

3. To configure L3-hand-off EPG under the ownership of shared tenant using the Shared VRF, run the following command:

```
efa tenant epg create --name <epg-name> --type l3-hand-off --tenant <shared-tenant-
name-owning-the-shared-vrf>
    --po <po-list> --switchport-mode <trunk|access> --ctag-range <ctag-range>
    --local-ip <ctag,device-ip:local-ip> --vrf <shared-vrf-owned-by-shared-tenant>
```

4. To configure BGP peer under the ownership of shared tenant using the shared VRF, run the following command:

```
efa tenant service bgp peer create --name <bgp-peer-name> --tenant <shared-tenant-
owning-the-shared-vrf>
    --ipv4-uc-nbr <device-ip,shared-vrf-owned-by-shared-tenant:bgp-peer-ip,bgp-peer-
remote-asn>
```

Shared VRF and Router Usecase with Examples

Learn the examples of various use cases of shared VRF, tenant, port channel, and router.

Topology

The following example shows the fabric topology:

```
efa fabric create --name fabric1 --type non-clos

efa fabric setting update --name fabric1
    --vni-auto-map No --backup-routing-enable Yes

efa fabric device add-bulk --name fabric1
    --rack rack1 --ip 10.20.246.25-26 --rack rack2 --ip 10.20.246.17-18
    --border-leaf-rack rack3 --border-leaf-ip 10.20.246.15-16
    --username admin --password password

efa fabric configure --name fabric1

efa fabric show

Fabric Name: default, Fabric Description: Default Fabric, Fabric Stage: 3, Fabric Type:
clos, Fabric Status: created
+--------+-----+-----+-----+------+-------+------+-----------+--------+------+---+
| IP     | POD | HOST| ASN | ROLE | DEVICE| APP  | CONFIG    | PENDING| VTLB | LB|
| ADDRESS|     | NAME|     |      | STATE | STATE| GEN REASON| CONFIGS| ID   | ID|
+--------+-----+-----+-----+------+-------+------+-----------+--------+------+---+

Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric Status:
settings-updated

Updated Fabric Settings: BGP-LL

+------------+---+------+-------+-------+-------------+------------+------------
+-------------------------+------+---+
```

```
| IP ADDRESS |POD| HOST | ASN    | ROLE   | DEVICE
STATE| APP STATE   | CONFIG GEN | PENDING CONFIGS          | VTLB |LB |
|            |   | NAME |        |        |             |
| REASON     |                          | ID    |ID |
+------------+---+------+--------+--------+-------------+------------+------------
+--------------------------+------+---+
| 10.20.246.1|   | SLX-1| 64512 | Spine | provisioned
| cfg in-sync| NA         | NA                       | NA    | 1 |
| 10.20.246.7|   | SLX  | 65000 | Leaf   | provisioning|
cfg ready  | IA,IU,MD,DA| SYSP-C,MCT-C,MCT-PA,     | 2     | 1 |
|            |   |      |        |        | failed      |
|            | BGP-C,INTIP-C,EVPN-C,O-C|      |   |
| 10.20.246.8|   | slx-8| 65000 | Leaf   | provisioned
| cfg in-sync| NA         | NA                       | 2     | 1 |
+------------+---+------+--------+--------+-------------+------------+------------
+--------------------------+------+---+
FABRIC SETTING:
BGPLL - BGP Dynamic Peer Listen Limit, BGP-MD5 - BGP MD5
Password , BFD-RX - Bfd Rx Timer, BFD-TX - Bfd Tx Timer, BFD-MULTIPLIER - Bfd multiplier,
BFD-ENABLE - Enable Bfd, BGP-MULTIHOP - BGP ebgp multihop,
P2PLR - Point-to-Point Link Range, MCTLR - MCT Link Range, LOIP - Loopback IP Range


CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/
Update, PLC/PLD/PLU - IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway
Delete/Update, EU/ED - Evpn Delete/Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd,
ASN - Asn Update, SYS - System Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update,
BGPLL - BGP Listen Limit, POU - Port Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP
- System Properties, INTIP - Interface IP, BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete
```

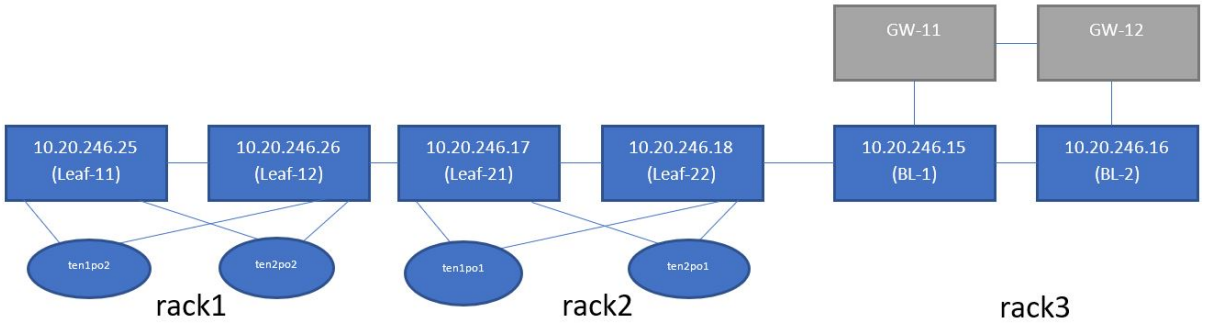

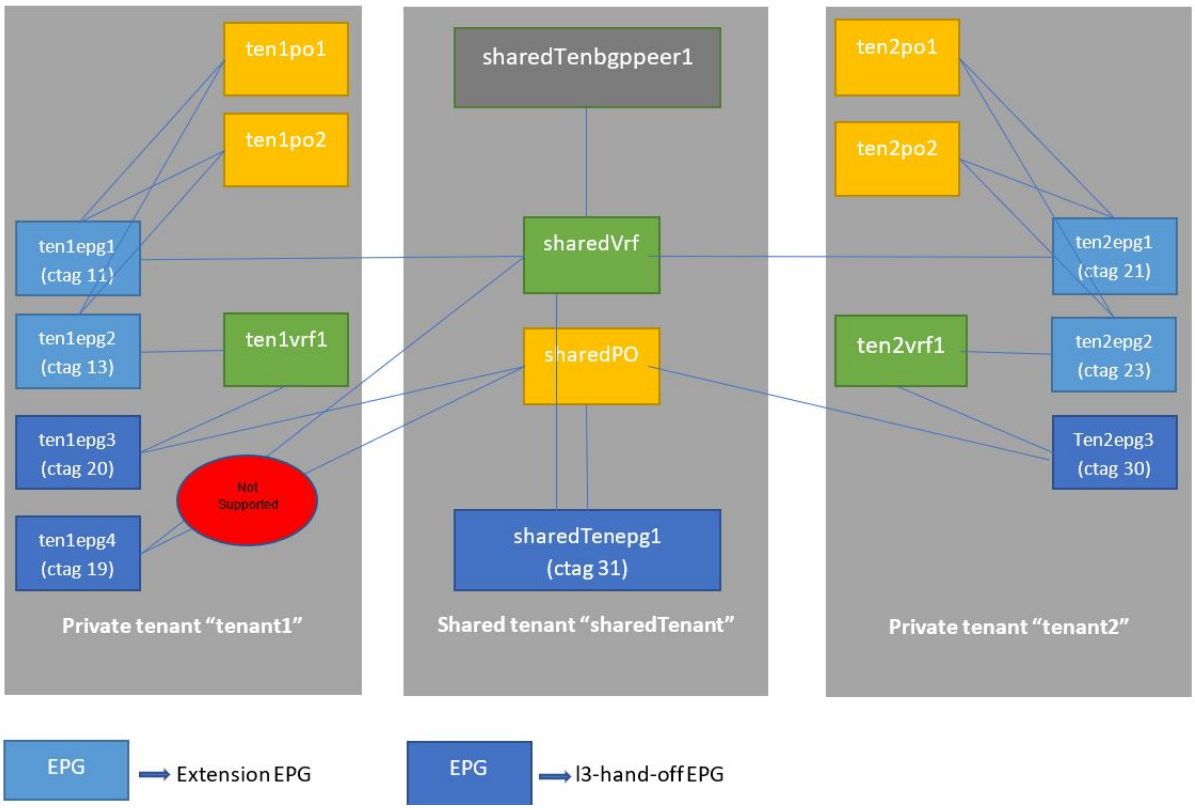

Figure 24: Shared VRF Configuration Overview

Shared Tenant and Private Tenant Configuration

The following example configures shared and private tenant:

```
efa tenant create --name sharedTenant --type shared --port
10.20.246.15[0/31],10.20.246.16[0/31]
    --vrf-count 10 --l3-vni-range 31001-31020
    --vlan-range 31-40 --l2-vni-range 30011-30020

efa tenant create --name tenant1 --port
10.20.246.17[0/11-20],10.20.246.18[0/11-20],10.20.246.25[0/11-20],10.20.246.26[0/11-20]
    --vlan-range 11-20 --l2-vni-range 10011-10020 --vrf-count 10 --l3-vni-range
11001-11020

efa tenant create --name tenant2 --port
10.20.246.17[0/21-30],10.20.246.18[0/21-30],10.20.246.25[0/21-30],10.20.246.26[0/21-30]
    --vlan-range 21-30 --l2-vni-range 20011-20020 --vrf-count 10 --l3-vni-range
21001-21020

efa tenant show
+------------+---------+-------+-------------+-------------+------+-------
+------------------------+
|    Name    |  Type   | VLAN  | L2VNI
Range | L3VNI Range | VRF   | Enable|         Ports         |
|            |         | Range |             |
| Count| BD    |                      |
+------------+---------+-------+-------------+-------------+------+-------
+------------------------+
|sharedTenant| shared  | 31-40 | 30011-30020 |
31001-31020 | 10    | false |  10.20.246.15[0/31]    |
|            |         |       |             |
```

```
|      |       |  10.20.246.16[0/31]   |
+------------+---------+-------+-------------+-------------+------+-------
+------------------------+
|  tenant1   | private | 11-20 | 10011-10020
| 11001-11020 | 10   | false | 10.20.246.18[0/11-20] |
|            |         |       |             |
|      |       | 10.20.246.17[0/11-20] |
|            |         |       |             |
|      |       | 10.20.246.25[0/11-20] |
|            |         |       |             |
|      |       | 10.20.246.26[0/11-20] |
+------------+---------+-------+-------------+-------------+------+-------
+------------------------+
|  tenant2   | private | 21-30 | 20011-20020
| 21001-21020 | 10   | false | 10.20.246.26[0/21-30] |
|            |         |       |             |
|      |       | 10.20.246.18[0/21-30] |
|            |         |       |             |
|      |       | 10.20.246.17[0/21-30] |
|            |         |       |             |
|      |       | 10.20.246.25[0/21-30] |
+------------+---------+-------+-------------+-------------+------+-------
+------------------------+
```

Shared PO and Private PO Configuration

The following example configures shared and private port channel:

```
efa tenant po create --name sharedPO --tenant sharedTenant
    --port 10.20.246.15[0/31],10.20.246.16[0/31] --speed 10Gbps --negotiation active

efa tenant po create --name ten1po1 --tenant tenant1
    --port 10.20.246.17[0/11],10.20.246.18[0/11] --speed 10Gbps --negotiation active

efa tenant po create --name ten1po2 --tenant tenant1
    --port 10.20.246.25[0/11],10.20.246.26[0/11] --speed 10Gbps --negotiation active
efa tenant po create --name ten2po1 --tenant tenant2
    --port 10.20.246.17[0/21],10.20.246.18[0/21] --speed 10Gbps --negotiation active
efa tenant po create --name ten2po2 --tenant tenant2
    --port 10.20.246.25[0/21],10.20.246.26[0/21] --speed 10Gbps --negotiation active
efa tenant po show
+----------+--------------+----+--------+-------------+----------+---------
+--------------------+------------+-------------+-------------+
|   Name   |    Tenant    | ID | Speed  | Negotiation
| Min Link |  Lacp   |       Ports        |   State    |  Dev State  |  App State  |
|          |         |    |        |            |
Count   | Timeout |                   |            |             |             |
+----------+--------------+----+--------+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| sharedPO | sharedTenant | 1  | 10Gbps |   active
|    1     |  long   | 10.20.246.16[0/31] | po-created | provisioned | cfg-in-sync |
|          |         |    |        |            |
|         | 10.20.246.15[0/31] |            |             |             |
+----------+--------------+----+--------+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| ten1po1  |   tenant1    | 1  | 10Gbps |   active
|    1     |  long   | 10.20.246.18[0/11] | po-created | provisioned | cfg-in-sync |
|          |         |    |        |            |
|         | 10.20.246.17[0/11] |            |             |             |
+----------+--------------+----+--------+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| ten1po2  |   tenant1    | 1  | 10Gbps |   active
|    1     |  long   | 10.20.246.25[0/11] | po-created | provisioned | cfg-in-sync |
|          |         |    |        |
```

```
|          |         | 10.20.246.26[0/11] |            |             |             |
+----------+---------+----+---------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
| ten2po1  |   tenant2    | 2  | 10Gbps |   active
|    1     |  long   | 10.20.246.18[0/21] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |
|          |         | 10.20.246.17[0/21] |            |             |             |
+----------+---------+----+---------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
| ten2po2  |   tenant2    | 2  | 10Gbps |   active
|    1     |  long   | 10.20.246.25[0/21] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |
|          |         | 10.20.246.26[0/21] |            |             |             |
+----------+---------+----+---------+-------------+-----------+---------
+--------------------+------------+-------------+-------------+
```

Shared VRF and Private VRF

The following example configures shared and private VRF:

```
efa tenant vrf create --name sharedVrf --tenant sharedTenant

efa tenant vrf create --name ten1vrf1 --tenant tenant1

efa tenant vrf create --name ten2vrf1 --tenant tenant2

efa tenant vrf show

+-----------+--------------+--------------+---------------------+--------------+----------
+-----------+-----------+------------+-----------------+-----------+
|   Name    |    Tenant    | Routing Type | Centralized Routers | Redistribute | Max Path
| Local Asn | Enable GR |   State    |    Dev State    | App State |
+-----------+--------------+--------------+---------------------+--------------+----------
+-----------+-----------+------------+-----------------+-----------+
| sharedVrf | sharedTenant | distributed  |                     |  connected   |    8
|           |   false   | vrf-create | not-provisioned | cfg-ready |
+-----------+--------------+--------------+---------------------+--------------+----------
+-----------+-----------+------------+-----------------+-----------+
| ten1vrf1  |   tenant1    | distributed  |                     |  connected   |    8
|           |   false   | vrf-create | not-provisioned | cfg-ready |
+-----------+--------------+--------------+---------------------+--------------+----------
+-----------+-----------+------------+-----------------+-----------+
| ten2vrf1  |   tenant2    | distributed  |                     |  connected   |    8
|           |   false   | vrf-create | not-provisioned | cfg-ready |
+-----------+--------------+--------------+---------------------+--------------+----------
+-----------+-----------+------------+-----------------+-----------+
```

Shared VRF: Inter POD Routing

The following example configures inter POD routing using shared VRF:

- Endpoint groups (EPGs) owned by different Private Tenants using the shared VRF

```
efa tenant epg create --name ten1epg1 --tenant tenant1 --po ten1po1,ten1po2 --
switchport-mode trunk
         --ctag-range 11 --anycast-ip 11:10.0.11.1/24 –vrf sharedVrf
efa tenant epg create --name ten2epg1 --tenant tenant2 --po ten2po1,ten2po2 --
switchport-mode trunk
        --ctag-range 21 --anycast-ip 21:10.0.21.1/24 --vrf sharedVrf
```

- Endpoint groups (EPGs) owned by different Private Tenants using their own private VRF:

```
efa tenant epg create --name ten1epg2 --tenant tenant1 --po ten1po1,ten1po2 --
switchport-mode trunk
         --ctag-range 13 --anycast-ip 13:10.0.13.1/24 --vrf ten1vrf1
```

```
efa tenant epg create --name ten2epg2 --tenant tenant2 --po ten2po1,ten2po2 --
switchport-mode trunk
        --ctag-range 23 --anycast-ip 23:10.0.23.1/24 --vrf ten2vrf1
```

```
efa tenant epg show
+----------+---------+-----------+-------+---------+------------+-------------
+------------+-----------+-------+-------+
|   Name   | Tenant  |   Type    | Ports |
PO    | SwitchPort | Native Vlan | Ctag Range |    Vrf    | L3Vni | State |
|          |         |           |       |         |    Mode
|   Tagging   |            |           |       |       |
+----------+---------+-----------+-------+---------+------------+-------------
+------------+-----------+-------+-------+
| ten1epg1 | tenant1 | extension |       | ten1po2
|   trunk    |    false    |     11     | sharedVrf | 31001 |       |
|          |         |           |       | ten1po1 |
|             |            |           |       |       |
+----------+---------+-----------+-------+---------+------------+-------------
+------------+-----------+-------+-------+
| ten1epg2 | tenant1 | extension |       |
ten1po1 |   trunk    |    false    |     13     | ten1vrf1  | 11001 |       |
|          |         |           |       | ten1po2 |
|             |            |           |       |       |
+----------+---------+-----------+-------+---------+------------+-------------
+------------+-----------+-------+-------+
| ten2epg1 | tenant2 | extension |       |
ten2po1 |   trunk    |    false    |     21     | sharedVrf | 31001 |       |
|          |         |           |       | ten2po2 |
|             |           |           |       |       |
|          |         |           |       |         |
|             |            |           |       |       |
+----------+---------+-----------+-------+---------+------------+-------------
+------------+-----------+-------+-------+
| ten2epg2 | tenant2 | extension |       |
ten2po2 |   trunk    |    false    |     23     | ten2vrf1  | 21001 |       |
|          |         |           |       | ten2po1 |
|             |            |           |       |       |
|          |         |           |       |         |
|             |            |           |       |       |
+----------+---------+-----------+-------+---------+------------+-------------
+------------+-----------+-------+-------+
```

L3 Hand-off using EPG and BGP Peer Owned by Shared Tenant Using Shared VRF

The following example configures an L3 Hand-off:

- Endpoint Group (EPG) owned by Shared Tenant handling off Shared VRF

```
efa tenant epg create --name sharedTenepg1 --tenant sharedTenant --type l3-hand-off
     --po sharedPO --switchport-mode trunk --ctag-range 31
     --vrf sharedVrf --local-ip 31,10.20.246.15:10.0.31.2/24 --local-ip
31,10.20.246.16:10.0.31.2/24
```

- BGP peer owned by Shared Tenant handling off Shared VRF

```
efa tenant service bgp peer create --name sharedTenbgppeer1 --tenant sharedTenant
       --ipv4-uc-nbr 10.20.246.15,sharedVrf:10.0.31.3,50000
       --ipv4-uc-nbr 10.20.246.16,sharedVrf:10.0.31.3,50000
```

```
efa tenant epg show
+---------------+--------------+-------------+-------+----------+------------
+-------------+------------+-----------+-------+-------+
|     Name      |    Tenant    |    Type     | Ports
|    PO    | SwitchPort | Native Vlan | Ctag Range |    Vrf    | L3Vni | State |
|               |              |             |       |
|    Mode    |   Tagging   |            |           |       |       |
```

```
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
| sharedTenepg1 | sharedTenant | l3-hand-off |
| sharedPO |   trunk    |    false    |     31     | sharedVrf | 31001 |       |
|               |            |             |            |
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten1epg2    |   tenant1    |  extension  |
| ten1po2  |   trunk    |    false    |     13     | ten1vrf1  | 11001 |       |
|               |            |             |            | ten1po1
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten1epg3    |   tenant1    | l3-hand-off |
| sharedPO |   trunk    |    false    |     20     | ten1vrf1  | 11001 |       |
|               |            |             |            |
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten1epg4    |   tenant1    | l3-hand-off |
| sharedPO |   trunk    |    false    |     19     | sharedVrf | 31001 |       |
|               |            |             |            |
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten1epg1    |   tenant1    |  extension  |
| ten1po2  |   trunk    |    false    |     11     | sharedVrf | 31001 |       |
|               |            |             |            |
ten1po1  |           |            |           |           |       |       |
|               |            |             |            |
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten2epg3    |   tenant2    | l3-hand-off
|       | sharedPO |   trunk    |    false    |     30     | ten2vrf1  | 21001 |       |
|               |            |             |            |
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten2epg1    |   tenant2    |  extension
|       | ten2po2  |   trunk    |    false    |     21     | sharedVrf | 31001 |       |
|               |            |             |            |
ten2po1  |           |            |           |           |       |       |
|               |            |             |            |
|           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
|   ten2epg2    |   tenant2    |  extension
|       | ten2po1  |   trunk    |    false    |     23     | ten2vrf1  | 21001 |       |
|               |            |             |            |
ten2po2  |           |            |           |           |       |       |
+--------------+-------------+-------------+-------+----------+------------
+-------------+-----------+-----------+-------+-------+
efa tenant service bgp peer show --detail
==========================================
Name             : sharedTenbgppeer1
Tenant           : sharedTenant
State            : bs-state-created
Description      :

Static Peer
-----------
        Device IP        : 10.20.246.15
```

```
        VRF              : sharedVrf
        AFI              : ipv4
        SAFI             : unicast
        Remote IP        : 10.0.31.3
        Remote ASN       : 50000
        Next Hop Self    : false
        Update Source IP :
        BFD Enabled      : false
        BFD Interval     : 0
        BFD Rx           : 0
        BFD Multiplier   : 0
        MD5 Password     :
        Dev State        : provisioned
        App State        : cfg-in-sync

        Device IP        : 10.20.246.16
        VRF              : sharedVrf
        AFI              : ipv4
        SAFI             : unicast
        Remote IP        : 10.0.31.3
        Remote ASN       : 50000
        Next Hop Self    : false
        Update Source IP :
        BFD Enabled      : false
        BFD Interval     : 0
        BFD Rx           : 0
        BFD Multiplier   : 0
        MD5 Password     :
        Dev State        : provisioned
        App State        : cfg-in-sync

Dynamic Peer
------------
        0 Records
=========================================
```

Sharing Multiple VRFs with the Same RT (route-target)

You can configure VRFs to share with the same route target.

About This Task

Follow this procedure to share multiple VRFs with the same route target.

Procedure

1. If you are running EFA 2.5.5 or above, run the following command:

```
efa tenant vrf create --tenant "ten1" --name "ten1vrf1" --routing-type "distributed"
--rt-type import --rt 100:100 --rt-type export --rt 100:100 --rt-type import --rt
200:200 --rt-type export --rt 200:200 --rt-type import --rt 300:300 --rt-type export
--rt 400:400 --max-path 8 --redistribute connected

efa tenant vrf create --tenant "ten2" --name "ten2vrf1" --routing-type "distributed"
--rt-type import --rt 100:100 --rt-type export --rt 100:100 --rt-type import --rt
200:200 --rt-type export --rt 200:200 --rt-type import --rt 300:300 --rt-type export
--rt 400:400 --max-path 8 --redistribute connected
```

2. If you are running EFA 2.5.5 or above, verify the switch configuration on SLX device.

```
L1# show running-config vrf
bvrf ten1vrf1
rd 172.31.254.69:1
evpn irb ve 8192
address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
!
vrf ten2vrf1
rd 172.31.254.69:2
evpn irb ve 8190
address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
!
```

```
L2# show running-config vrf
bvrf ten1vrf1
rd 172.31.254.70:1
evpn irb ve 8192
address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
!
vrf ten2vrf1
rd 172.31.254.70:2
evpn irb ve 8190
address-family ipv4 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
address-family ipv6 unicast
  route-target export 100:100 evpn
  route-target export 200:200 evpn
  route-target export 400:400 evpn
  route-target import 100:100 evpn
  route-target import 200:200 evpn
  route-target import 300:300 evpn
!
!
```

Configure Tenant Admin Access to Shared Tenant Resources or Entities

In XCO versions prior to 3.0.0, running the REST GET API or the equivalent CLI without `tenant` filter disables the tenant admin to view the resources or entities owned by the tenant admin and the resources or entities owned by the shared tenant.

In XCO versions 3.0.0 or above, running the REST GET API or the equivalent CLI without `tenant` filter enables the tenant admin to view the resources or entities owned by the tenant admin and the resources or entities owned by the shared tenant.

About This Task

Follow this procedure to configure tenant admin access to shared tenant resources.

Procedure

1. Log in to XCO as a root user.

```
(efa:root)root@administrator-00:~# efa tenant show
+-------------+--------+-------+-------------+-------------+------+-------
+-------------------+------------------+
|    Name     | Type   | VLAN  | L2VNI Range
| L3VNI Range | VRF  | Enable|       Ports       | Mirror           |
|             |        | Range |             |
| Count| BD    |                   | Destination Ports|
+-------------+--------+-------+-------------+-------------+------+-------
+-------------------+------------------+
| roottenant  |private | 2-20  | 10000-10099
| 10110-10119 | 10   | false | 10.20.246.4[0/20] | 10.20.246.4[0/21]|
|             |        |       |             |
|      |       | 10.20.246.3[0/20] | 10.20.246.3[0/21]|
+-------------+--------+-------+-------------+-------------+------+-------
+-------------------+------------------+
|sharedtenant |shared  | 2-20  | 20000-20099 |
20110-20119 | 10   | false | 10.20.246.4[0/22] | 10.20.246.3[0/23]|
|             |        |       |             |
|      |       | 10.20.246.3[0/22] | 10.20.246.4[0/23]|
+-------------+--------+-------+-------------+-------------+------+-------
+-------------------+------------------+
|     t1      |private | 2-20  | 30000-30099
| 30110-30119 | 10   | false | 10.20.246.4[0/24] | 10.20.246.3[0/25]|
|             |        |       |             |
|      |       | 10.20.246.3[0/24] | 10.20.246.4[0/25]|
+-------------+--------+-------+-------------+-------------+------+-------
+-------------------+------------------+
|     t2      |private | 2-20  | 40000-40099
| 40110-40119 | 10   | false | 10.20.246.4[0/26] | 10.20.246.4[0/27]|
|             |        |       |             |
|      |       | 10.20.246.3[0/26] | 10.20.246.3[0/27]|
+-------------+--------+-------+-------------+-------------+------+-------
+-------------------+------------------+
Tenant Details

(efa:root)root@administrator-00:~# efa tenant po show
+----------+--------------+----+---------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
|   Name   |    Tenant    | ID | Speed   | MTU | Negotiation
| Min Link |  Lacp   |       Ports       |   State    |  Dev State  |  App State  |
|          |              |    |         |     |
|  Count   | Timeout |                   |            |             |             |
+----------+--------------+----+---------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
|  rootpo  |  roottenant  | 2  | 10Gbps |     |   active
|    1     |  long   | 10.20.246.4[0/20] | po-created | provisioned | cfg-in-sync |
|          |              |    |         |     |
|          |         | 10.20.246.3[0/20] |            |             |             |
+----------+--------------+----+---------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
| sharedpo | sharedtenant | 3  | 10Gbps |     |   active
|    1     |  long   | 10.20.246.4[0/22] | po-created | provisioned | cfg-in-sync |
|          |              |    |         |     |
|          |         | 10.20.246.3[0/22] |            |             |             |
+----------+--------------+----+---------+-----+-------------+----------+---------
+--------------------+------------+-------------+-------------+
|   po1    |      t1      | 4  | 10Gbps |     |   active
|    1     |  long   | 10.20.246.4[0/24] | po-created | provisioned | cfg-in-sync |
|          |              |    |         |     |
|          |         | 10.20.246.3[0/24] |            |             |             |
+----------+--------------+----+---------+-----+-------------+----------+---------
```

```
+------------------+-----------+-------------+-------------+
|   po2    |      t2      | 5  | 10Gbps |     |   active
|    1     |  long   | 10.20.246.4[0/26] | po-created | provisioned | cfg-in-sync |
|          |             |    |        |     |
|          |         | 10.20.246.3[0/26] |            |             |             |
+----------+-------------+----+--------+-----+-------------+----------+---------
+------------------+-----------+-------------+-------------+
Port Channel Details


(efa:root)root@administrator-00:~# efa tenant vrf show
+-----------+-------------+-------------+-----------+-------------+-----+------
+-------+--------------------+---------------+------------+
|   Name    |    Tenant    | Routing Type | Centralized|
Redistribute | Max | Local| Enable|       State        | Dev State      | App State  |
|           |              |              | Routers    |
| Path| Asn  | GR    |                    |                |            |
+-----------+-------------+-------------+-----------+-------------+-----+------
+-------+--------------------+---------------+------------+
|  rootvrf  |  roottenant  | distributed  |            |
connected   |  8  |      | false | vrf-device-created |  provisioned   |cfg-in-sync |
+-----------+-------------+-------------+-----------+-------------+-----+------
+-------+--------------------+---------------+------------+
| sharedvrf | sharedtenant | distributed  |            |
connected   |  8  |      | false | vrf-device-created |  provisioned   |cfg-in-sync |
+-----------+-------------+-------------+-----------+-------------+-----+------
+-------+--------------------+---------------+------------+
|   myv1    |      t1      | distributed  |            |
connected   |  8  |      | false | vrf-device-created |  provisioned   |cfg-in-sync |
+-----------+-------------+-------------+-----------+-------------+-----+------
+-------+--------------------+---------------+------------+
|   myv2    |      t2      | distributed  |            |
connected   |  8  |      | false |    vrf-created     |not-provisioned | cfg-ready  |
+-----------+-------------+-------------+-----------+-------------+-----+------
+-------+--------------------+---------------+------------+
Vrf Details

(efa:root)root@administrator-00:~# efa tenant epg show
+-----------+-------------+-------------+-------+--------+-----------+-------------
+------+----------+-------+---------------------+
|   Name    |    Tenant    |    Type     | Ports
|   PO   | SwitchPort | Native Vlan | Ctag |   Vrf    | L3Vni |      State          |
|          |              |             |       |
|    Mode    |   Tagging   | Range|          |       |                     |
+-----------+-------------+-------------+-------+--------+-----------+-------------
+------+----------+-------+---------------------+
|  rootepg  |  roottenant  |  extension  |
| rootpo |   trunk    |    false    |  10  | rootvrf  | 10111 | epg-with-port-group |
|          |             |             |       |
|            |             |      |          |       | -and-ctag-range     |
+-----------+-------------+-------------+-------+--------+-----------+-------------
+------+----------+-------+---------------------+
| sharedepg | sharedtenant | l3-hand-off |
|        |            |    false    |      |          |       |     epg-empty       |
|          |             |             |       |
|            |             |      |          |       |                     |
+-----------+-------------+-------------+-------+--------+-----------+-------------
+------+----------+-------+---------------------+
|   epg1    |      t1      |  extension  |
|  po1   |   trunk    |    false    |  11  |  myv1    | 30111 | epg-with-port-group |
|          |             |             |       |
|            |             |      |          |       | -and-ctag-range     |
+-----------+-------------+-------------+-------+--------+-----------+-------------
+------+----------+-------+---------------------+
```

```
|   epg2    |      t2      |  extension  |
|  po2   |   trunk    |    false    |  12  |sharedvrf | 20111 | epg-with-port-group |
|           |              |             |       |
|            |             |      |          |       | -and-ctag-range     |
+-----------+--------------+-------------+-------+-------+-------------+-------------
+------+----------+-------+---------------------+
EndpointGroup Details
```

2. Log in to XCO as a tenant user.

```
(efa:t1user)root@administrator-00:~# efa tenant show
+--------------+---------+------------+-------------+-------------+-----------
+-----------+-------------------+--------------------------+
|     Name     |  Type   | VLAN Range | L2VNI Range
| L3VNI Range | VRF Count | Enable BD |       Ports       | Mirror Destination Ports |
+--------------+---------+------------+-------------+-------------+-----------
+-----------+-------------------+--------------------------+
| sharedtenant | shared  |    2-20    | 20000-20099
| 20110-20119 |    10     |   false   | 10.20.246.4[0/22] |    10.20.246.4[0/23]     |
|               |         |            |             |
|           |           | 10.20.246.3[0/22] |    10.20.246.3[0/23]     |
+--------------+---------+------------+-------------+-------------+-----------
+-----------+-------------------+--------------------------+
|      t1      | private |    2-20    | 30000-30099
| 30110-30119 |    10     |   false   | 10.20.246.4[0/24] |    10.20.246.4[0/25]     |
|               |         |            |             |
|           |           | 10.20.246.3[0/24] |    10.20.246.3[0/25]     |
+--------------+---------+------------+-------------+-------------+-----------
+-----------+-------------------+--------------------------+
Tenant Details

(efa:t1user)root@administrator-00:~# efa tenant po show
+----------+--------------+----+--------+-----+-------------+----------+---------
+-------------------+------------+-------------+-------------+
|   Name   |    Tenant    | ID | Speed  | MTU | Negotiation
| Min Link |  Lacp   |       Ports       |   State    |  Dev State  |  App State  |
|          |              |    |        |     |
|  Count   | Timeout |                   |            |             |             |
+----------+--------------+----+--------+-----+-------------+----------+---------
+-------------------+------------+-------------+-------------+
| sharedpo | sharedtenant | 3  | 10Gbps |     |   active
|    1     |  long   | 10.20.246.4[0/22] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |     |
|          |         | 10.20.246.3[0/22] |            |             |             |
+----------+--------------+----+--------+-----+-------------+----------+---------
+-------------------+------------+-------------+-------------+
|   po1    |      t1      | 4  | 10Gbps |     |   active
|    1     |  long   | 10.20.246.4[0/24] | po-created | provisioned | cfg-in-sync |
|          |              |    |        |     |
|          |         | 10.20.246.3[0/24] |            |             |             |
+----------+--------------+----+--------+-----+-------------+----------+---------
+-------------------+------------+-------------+-------------+
Port Channel Details
(efa:t1user)root@administrator-00:~# efa tenant vrf show
+-----------+--------------+--------------+------------+--------------+-----+------
+-------+--------------------+-------------+-------------+
|   Name    |    Tenant    | Routing Type | Centralized|
Redistribute | Max | Local| Enable|       State        |  Dev State  |  App State  |
|           |              |              | Routers    |
| Path| Asn  | GR    |                   |             |             |
+-----------+--------------+--------------+------------+--------------+-----+------
+-------+--------------------+-------------+-------------+
| sharedvrf | sharedtenant | distributed  |
|  connected   |  8  |      | false | vrf-device-created | provisioned | cfg-in-sync |
+-----------+--------------+--------------+------------+--------------+-----+------
```

```
+-------+--------------------+-------------+-------------+
|   myv1    |      t1      | distributed  |
|  connected   |  8  |      | false | vrf-device-created | provisioned | cfg-in-sync |
+-----------+--------------+-------------+------------+--------------+-----+------
+-------+--------------------+-------------+-------------+
Vrf Details


(efa:t1user)root@administrator-00:~# efa tenant epg show
+----------+-------------+------------+------+-----+------------+------------+------
+-----+-------+----------------------+
|   Name   |    Tenant   |    Type    | Ports|
PO  | SwitchPort |Native Vlan | Ctag |Vrf  | L3Vni |      State          |
|          |             |            |      |     |    Mode
|  Tagging   | Range|     |       |                     |
+----------+-------------+------------+------+-----+------------+------------+------
+-----+-------+----------------------+
| sharedepg| sharedtenant| l3-hand-off|      |     |
|   false    |      |     |       |     epg-empty       |
|          |             |            |      |     |
|            |      |     |       |                     |
+----------+-------------+------------+------+-----+------------+------------+------
+-----+-------+----------------------+
|   epg1   |      t1     |  extension |      |
po1 |   trunk    |   false    | 11   |myv1 | 30111 | epg-with-port-group |
|          |             |            |      |     |
|            |      |     |       | -and-ctag-range     |
+----------+-------------+------------+------+-----+------------+------------+------
+-----+-------+----------------------+
EndpointGroup Details
```

Administered Partial Success

For a two-leaf MCT pair, you can enable configurations to succeed when only one node of the pair is reachable.

Overview

By default, when a REST operation succeeds on one device but fails on another, configuration changes are rolled back for both devices. For more information, see Rollback Scenarios for Data Consistency on page 72.

However, for a two-leaf MCT pair, you can administratively change the process to permit configuration to succeed even when one device is down. This process, called an administered partial success, is as follows.

- You use the **`efa inventory admin-state`** command to change the state of the unreachable device to "admin down." The device then goes into maintenance mode. For more information about changing a device state, see Administratively Manage a Device State on page 365.
- XCO filters out configurations destined for MCT pair as follows.
  - Create-related and delete-related configurations destined for the "admin up" device succeed.
  - Create-related configurations are not attempted for the "admin down" device, but the configurations are considered a success. These configurations are marked as pending, to be pushed to the device when it comes back up.

    - Delete-related configurations (de-configurations) are not attempted for the "admin down" device and the operation fails with an error in the REST response. You can retry these de-configurations after the device transitions to “admin up” state.

      XCO does not want to leave stale configurations on the devices because if stale configurations are left on the devices, then bringing the devices (with stale configurations) back into XCO are erroneous considering the full brownfield support is missing in XCO.
- When the device is again reachable, you change the state of the device to "admin-up."
- XCO pushes the pending configurations to the device, and the drift and reconcile process ensures that the configurations in XCO and the device are synchronized. For more information, see Drift and Reconcile on page 60.
- The device comes out of maintenance mode.

Tips and considerations

- You can use Switch Health Management to verify the reachability of a device. Use the **`--health-check-interval`** and **`--health-check-heartbeat-miss-threshold`** settings of the **`efa inventory device setting update`** command. For more information, see Monitor Device Health on page 459.
- You can retry the same CLI or REST operation after the “admin down” devices transition to “admin up” state so that the deconfiguration is attempted on all the devices. You can use the “force” option available in the REST API to forcefully delete the entities from XCO even in case of partial success topology.
- You can use the **`efa tenant debug device drift`** command to determine any drift between the intended XCO configuration and the device configuration. These commands also identify the app state and the dev state: **`efa tenant epg show`** and **`efa tenant po show`**.
- XCO blocks the tenant reconciliation API, and rest of the tenant APIs support partial success behavior.
- If a high-availability failover or restart occurs while a device is in "admin down" or "admin up" state, you must reapply the state.
- If an operation such as drift and reconcile or a firmware download is in progress when you submit the command to change the state, the command is blocked until the operation is complete.
- This feature is supported only for devices in an MCT pair. Standalone devices are not supported.
- You can change the status of only one device in an MCT pair to "admin down" to benefit from administered partial success.
    - When both devices are in "admin down" state, the topology is considered a complete failure. Configuration attempts on these devices are rejected and error messages are returned in the REST responses. Administered partial success is not applicable.

- When both devices are in "admin up" state, the topology is considered a complete success. Configuration attempts on these devices are accepted. Administered partial success is not applicable.

Behavior changes during "admin down" state

After a device state changes to "admin down," the following behavior changes occur.

- Switch Health Management does not trigger the drift and reconcile process.
- A device going into maintenance mode does not trigger the drift and reconcile process.
- The following commands are blocked from affecting the device.

Table 18: Blocked commands

Command type and name
Inventory commands
efa inventory device compare --ip
efa inventory drift-reconcile --ip
efa inventory device setting update --ip
efa inventory rma --ip
efa inventory config-backup execute --ip
efa inventory config-replay execute --ip
efa inventory device update --fabric
efa inventory device firmware-download prepare add --ip
efa inventory device update --ip
efa inventory device interface set-speed --ip
efa inventory device interface set-breakout --ip
efa inventory device interface unset-breakout --ip
efa inventory device interface set-mtu --ip
efa inventory device interface set-admin-state --ip
efa inventory device running-config persist --ip
Fabric commands
efa fabric configure --name
efa fabric device remove --ip <> --name <> Allowed with the --no-device-cleanup option.
efa fabric show-config --name
efa fabric topology show underlay --name
efa fabric topology show overlay --name
efa fabric topology show physical --name

Behavior changes during "admin up" state

After a device is returned to "admin up" state and after the drift and reconcile process is complete (which the state change triggers), Switch Health Management and drift and reconcile resume normal behavior. Also, the blocked commands are unblocked.

Administratively Manage a Device State

You can administratively manage the state of an SLX device using the XCO command line.

About This Task

You can change a state to up or down, delete a state from the history, and view the state history and the current state. For details about the command and its parameters, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

1. To change a device to the up state, run the following command:

```
$ efa inventory admin-state up --ip <device IP>
AdminStateUp [success]
Admin State Up execution UUID: 8d9fa0cf-dc76-42cc-ac7a-57902a47c1b2
```

   This example changes the state of a specified IP address and generates a UUID, which you can use in the **`efa inventory admin-state detail`** command:

2. To change a device to the down state, run the following command:

```
$ efa inventory admin-state down --ip <device IP>
AdminStateDown [success]
Admin State Down execution UUID: 28eb0845-7a7a-4851-b453-b3020c6900f2
```

   This example changes the state of a specified IP address and generates a UUID, which you can use in the **`efa inventory admin-state detail`** command.

3. To view the details of a state change, run the following command:

```
$ efa inventory admin-state detail --uuid 28eb0845-7a7a-4851-b453-b3020c6900f2
```

4. To view the history of the admin changes of a specified device, run the following command:

```
# efa inventory admin-state history --ip <device IP>
```

5. To display the admin state and the health check state of a device, run the following command:

```
$ efa inventory admin-state show --ip <device IP>
```

6. To delete the instance of the admin state change of a device, run the following command:

```
$ efa inventory admin-state delete --key <device IP or UUID>
```

Traffic Mirroring Overview

XCO supports traffic monitoring on both Clos and non-Clos (small data center) fabrics for troubleshooting issues with applications and fabrics. XCO performs traffic monitoring by means of packet mirroring in a cloud-native infrastructure solution and network functions virtualization in infrastructure deployments.

You can mirror the ingress and egress traffic from the following ports:

- Leaf ports connecting to the compute devices
- Leaf ports connecting to the neighboring MCT leaf device (ICL ports)
- Border leaf ports connecting to the external gateway
- Border leaf ports connecting to the neighboring MCT border leaf device (ICL ports)
- Spine ports connecting to leaf devices (Fabric non-ICL ports)
- Spine ports connecting to super-spine devices (Fabric non-ICL ports)
- Super-Spine ports connecting to spine and border-leaf devices (Fabric non-ICL ports)

There are two types of traffic mirroring:

1. In-band traffic mirroring
2. Out-of-band traffic mirroring

The following table describes the comparison between In-band and Out-of-band traffic mirroring solution:

In-band Mirroring	Out-of-band Mirroring
No additional hardware or ports	One additional switch, one reserved port on all leaf and border leaf switches
All configuration by XCO, no separate devices to be managed	Separate configuration on mirror switch through OOB mechanisms
All ingress information, including test access point (TAP) and VLAN, can be retained and used for classification	Ingress port information and possibly VLAN information, is not retained
Fabric needs to be measured for expected extra mirror traffic	Mirroring traffic has minimal impact on normal traffic and fabric capacity, no extra measurement needed
All functionality needs to be present in ingress leaf top of rack (ToR) switch	Minimal configuration needed on XCO, and dataplane support needed in the fabric
Extra tunnel configuration in fabric underlay	Fabric underlay is unmodified
Configuration of underlay tunnels to sink app breaks underlay/overlay separation	Tunnels to sink apps are outside the domain of fabric, and do not overlap
Cannot be applied for control port mirroring	Partial reuse possible for a common mirroring solution also on control network
Fabric has to be programmed for creating additional headers and remote destination reachability, underlay or overlay separation is lost	No fabric dependency on final encapsulation and forwarding toward sink
Egress ACL rule support minimal	Two level filtering possible, once in ingress switch, and once in the dedicated mirror switch, More complicated mirror rules can be cascaded.

In-band Mirroring	Out-of-band Mirroring
QoS support needed on tenant and mirrored traffic streams because they share the same fabric links	No QoS support needed, because links are separate
Cannot be leveraged for troubleshooting fabric issues due to reliance on fabric	Can be leveraged for troubleshooting fabric issues
Fabric admin needs to do all configuration because underlay routing modifications are needed	XCO tenant admin can create TAP sessions on the fabric switches, with pre-provisioning and custom provisioning of the configuration on mirror switch by fabric admin.

Note
For information about commands and supported parameters to configure traffic mirroring, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

In-band Traffic Mirroring

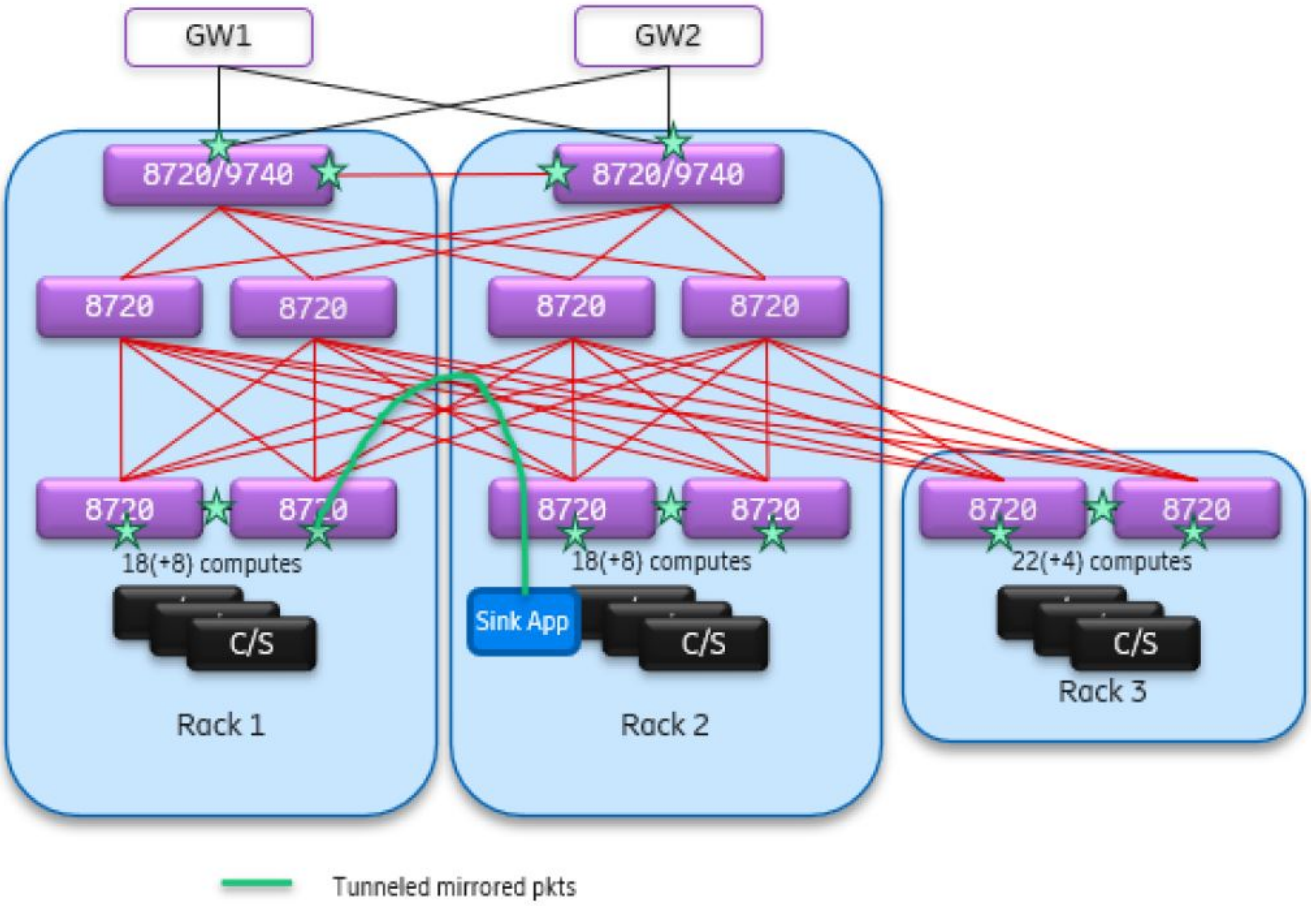

Figure 25: In-band traffic mirroring topology

- Fabric links are used for carrying mirrored and tenant traffic. Mirrored traffic needs to be encapsulated in ERSPAN headers to support multiple sessions.
- Separate tunnels are created over fabric links between the ingress leaf switch and either egress leaf switch or directly to the sink.
- Sink can be deployed in computes, in separate stand-alone servers, or outside the datacenter.
- No separate ports or devices are required for mirroring.

Out-of-band Traffic Mirroring

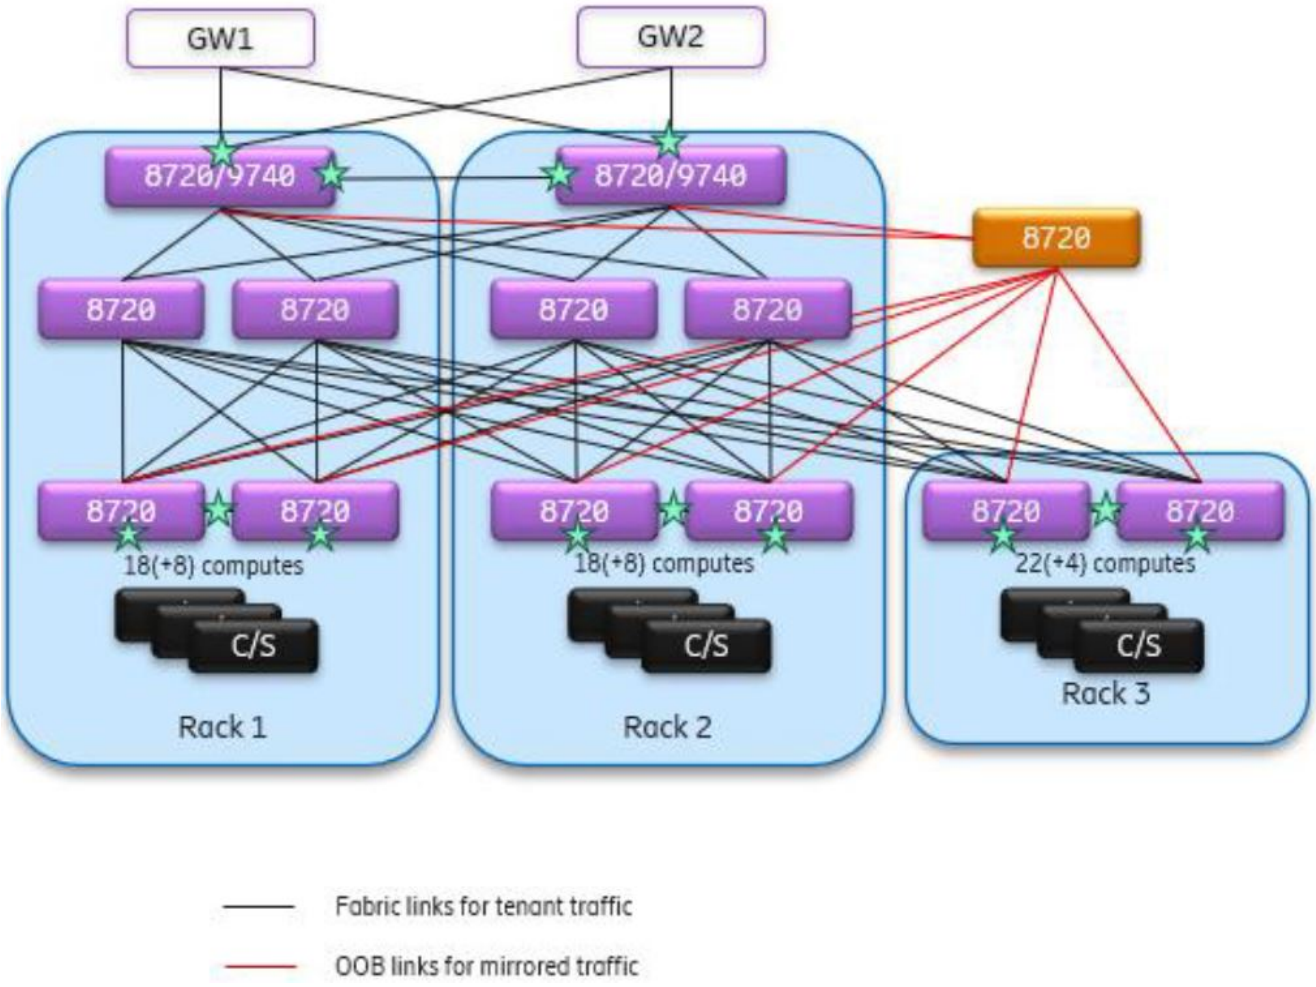

Figure 26: Out-of-band traffic mirroring topology

- The mirrored traffic is captured on the ingress or egress leaf switch and carried on a separate set of links to a separate add-on mirror switch.
- One port is reserved on each fabric leaf and border leaf switch, and connected to the mirror switch through separate OOB cabling.
- XCO configures basic mirroring sessions and actions on fabric switches.
- Advanced configuration on the mirror switch is handled separately, not by XCO.
- Connectivity between the TAP (traffic access point) sink (a class or function designed to receive incoming events from another object or function.) and the mirror switch can be configured and customized separately through OOB mechanism.
- Demultiplexing on sessions involving traffic access points (TAP) can be performed using filtering on packet header fields in the sink application.
- 100Gbps links are required between the mirror switch and each fabric switch.
- The mirror switch may be an 8720 or a specialized packet broker with advanced functions.

There are three types of out-of-band traffic mirroring:

- Port-based traffic mirroring
- Flow-based traffic mirroring
- VLAN-based traffic mirroring
- ICL port mirroring

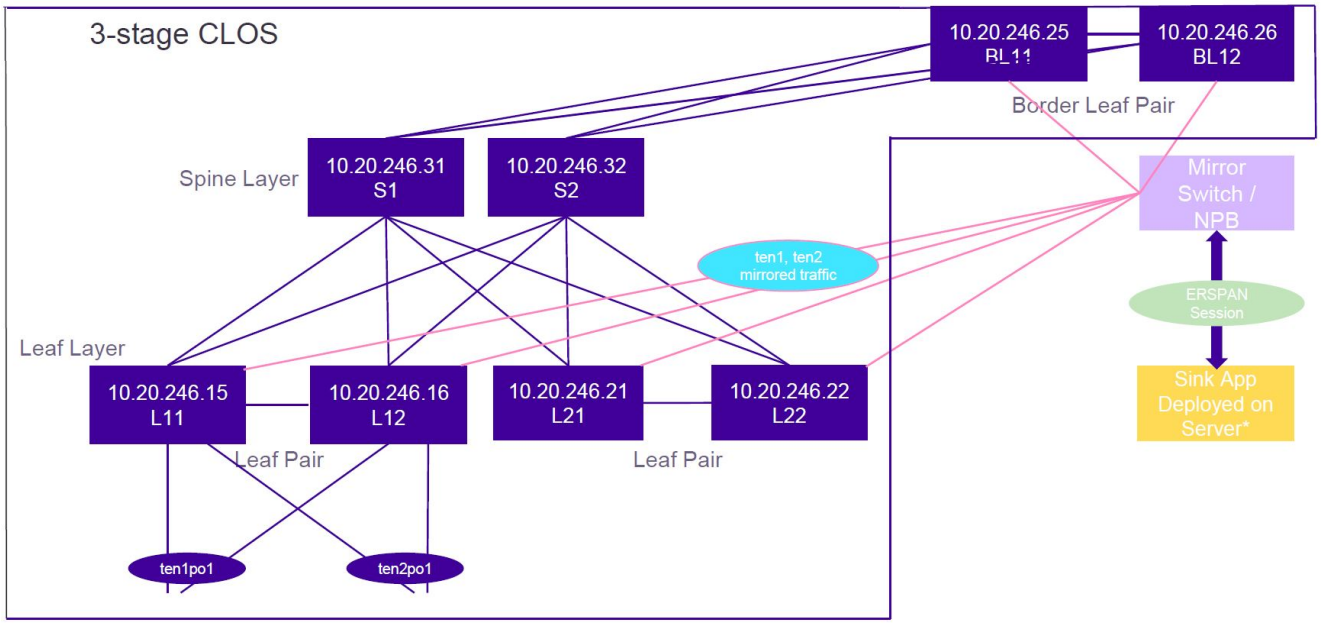

Figure 27: Port-based traffic mirroring topology

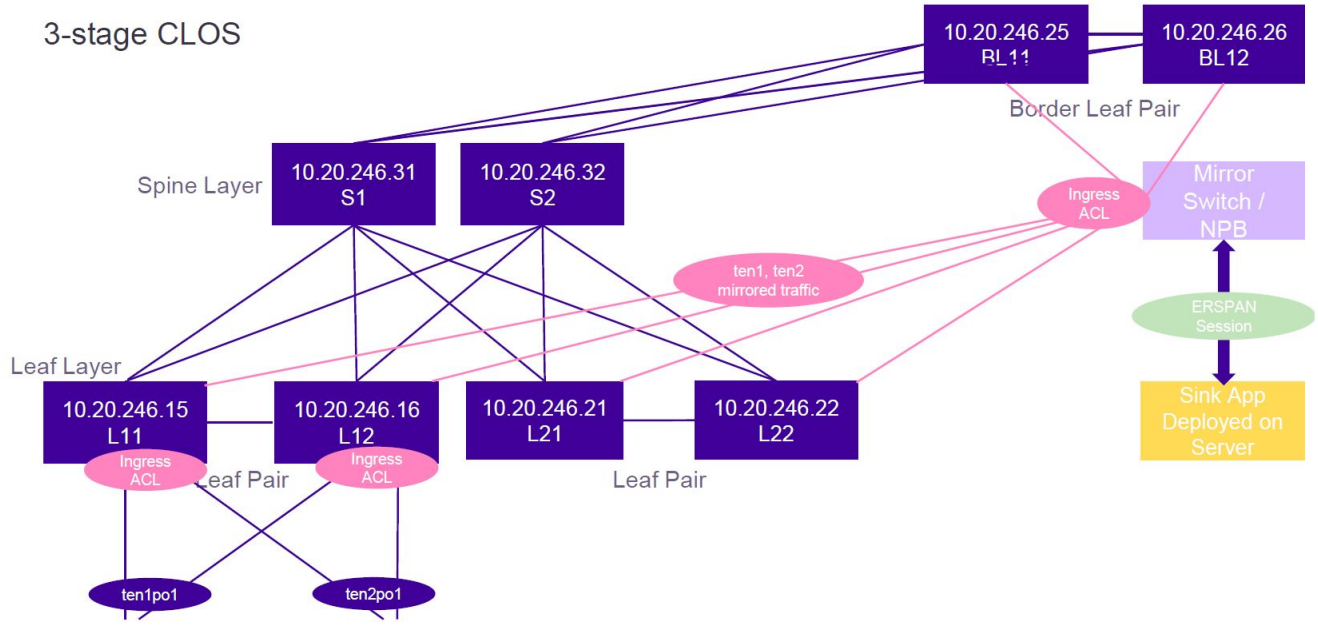

Figure 28: Flow-based traffic mirroring topology

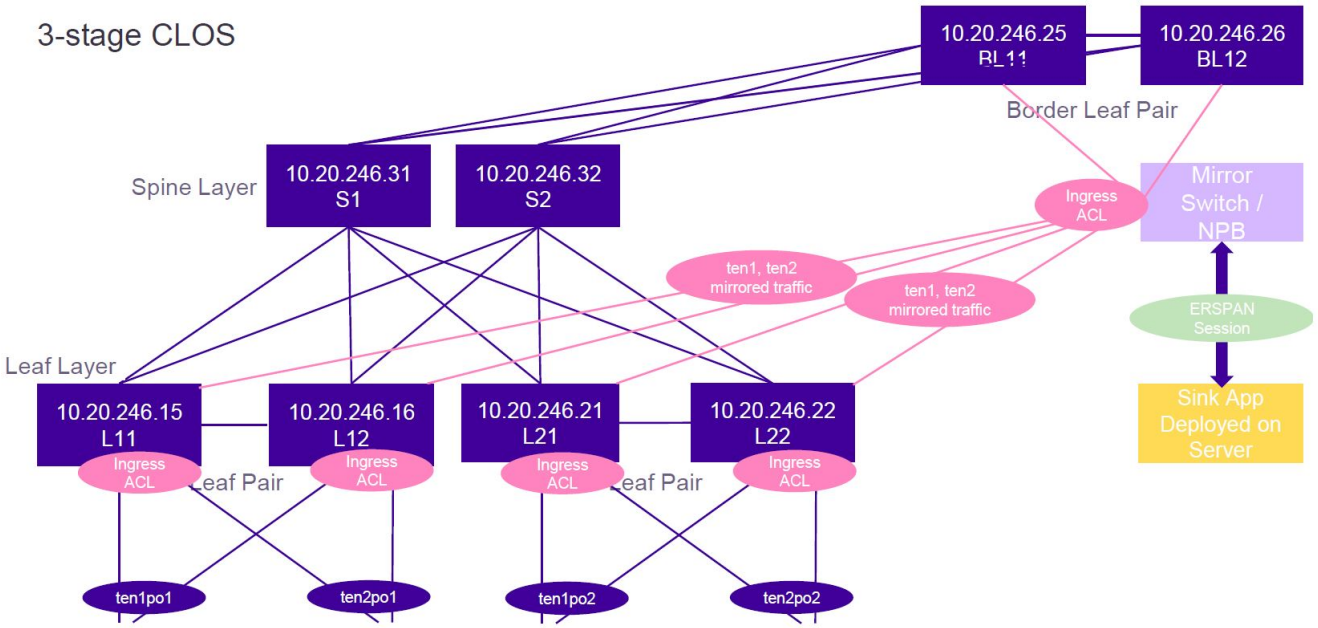

Figure 29: VLAN-based traffic mirroring topology

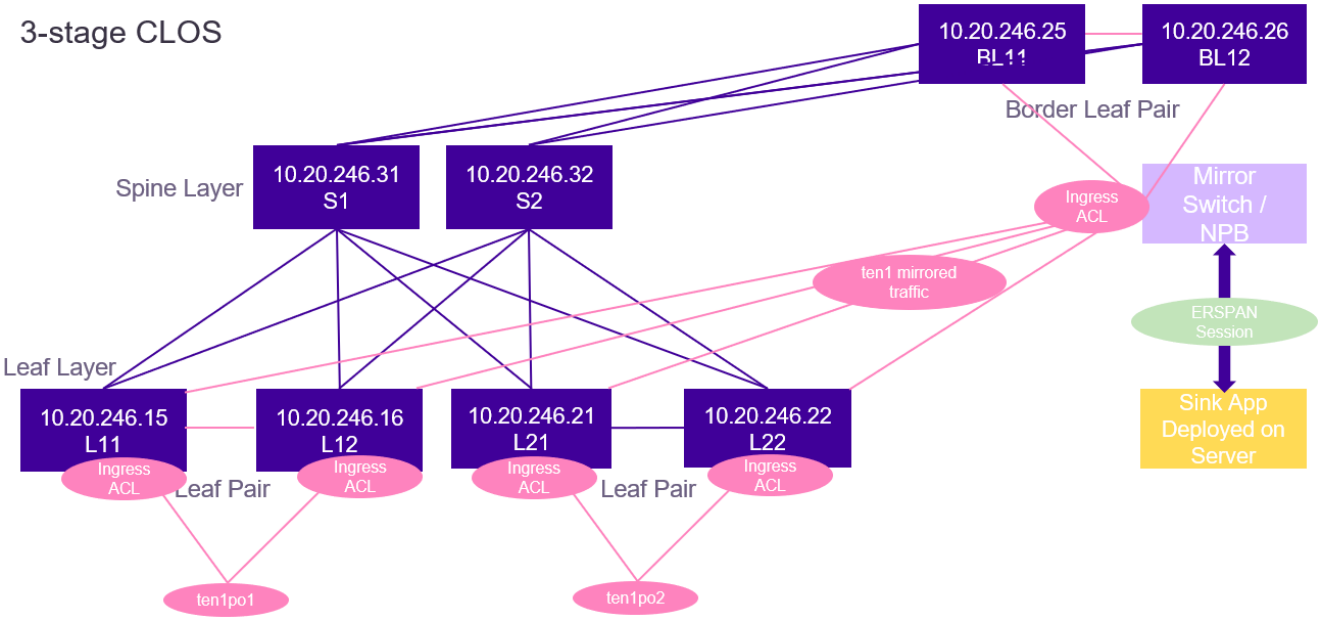

Figure 30: ICL port traffic mirroring topology

Support Matrix

Mirror Type	Mirror Destination Type	Mirror Source	Mirror Destination
Port-Based	Span	Ethernet	Ethernet
Port-Based	Span	Port-Channel	Ethernet

Mirror Type	Mirror Destination Type	Mirror Source	Mirror Destination
Flow-Based	Span	Ethernet	Ethernet
Flow-Based	Span	Port-Channel	Ethernet
Flow-Based	Span	Local VLAN (With device-ip)	Ethernet
Flow-Based	Span	Global VLAN (Without device-ip)	Ethernet
Flow-Based	Span	Local VLAN Range (Supported Range format a-b)	Ethernet
Flow-Based	Span	Global VLAN Range (Supported Range format a-b)	Ethernet

	np-mac-acl-in	np-mac-acl-out	np-ip-acl-in	np-ip-acl-out	np-ipv6-acl-in
L2 VLAN EPG	Supported	Supported	Not Supported	Not Supported	Not Supported
L3 VLAN EPG	Not Supported	Not Supported	Supported	Supported	Supported
L2 BD EPG	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
L3 BD EPG	Not Supported	Not Supported	Supported	Supported	Supported

	pp-mac-acl-in	pp-mac-acl-out	pp-ip-acl-in	pp-ip-acl-out	pp-ipv6-acl-in
Ethernet	Supported	Supported	Supported	Supported	Supported
Port-Channel	Supported	Not supported, except 9740	Supported	Supported	Not supported, except 9740

Note

- Only SPAN is supported as Destination Type and the Mirror Direction support is platform dependent.
- Only Ethernet interface is supported as Mirror Destination.
- Mirror Destination port value is mandatory for a mirror source when you create a mirror session, except for a Global VLAN SPAN.

Provision a Traffic Mirror Session

You can configure a traffic mirror session.

About This Task

Complete the following tasks to configure a traffic mirror session in your XCO fabric:

Procedure

1. Configure Port-Based Mirroring in a Multi-Tenant Architecture on page 373
2. Configure Flow-Based Mirroring in a Multi-Tenant Architecture on page 375
3. Configure VLAN-Based Mirroring in a Multi-Tenant Architecture on page 379
4. Configure ICL Port Mirroring in a Multi-Tenant Architecture on page 383
5. Configure Fabric Non-ICL Ports as Mirror Source on page 387

Configure Port-Based Mirroring in a Multi-Tenant Architecture

You can configure port-based mirroring in a multi-tenant architecture.

About This Task

Follow this procedure to configure post-based mirroring.

Procedure

1. Run the following commands to configure access control list applications on Ethernet or port channel and VLAN or virtual Ethernet:

```
efa tenant service mirror session create --name <session-name> --tenant <tenant-name>

   --source {<device-ip>,<eth | po | vlan>,<if-name>}
   --type {<source-device-ip>,<eth | po | vlan>,<source-if-name>:<port-based | flow-based>}

   --destination {<source-device-ip>,<eth | po | vlan>,<source-if-name> :
                <destination-device-ip>,<eth | po | vlan>,<destination-if-name}
   --destination-type {<source-device-ip>,< eth | po | vlan>,<source-if-name>:<span>}


   --direction {<source-device-ip>,< eth | po | vlan>,<source-if-name> : <tx | rx | both>}


(efa:root)root@node-2:~# efa tenant show
+------------+--------+------+-----+------+------+-------+----------------------+--------------------+
|    Name    |  Type  | VLAN |L2VNI| L3VNI| VRF  | Enable|         Ports        |   Mirroring Ports  |
|            |        | Range|Range| Range| Count| BD    |                      |                    |
+------------+--------+------+-----+------+------+-------+----------------------+--------------------+
|sharedTenant| shared |      |     |      | 0    | false |                      | 10.20.246.15[0/31]|
|            |        |      |     |      |      |       |                      | 10.20.246.16[0/31]|
|            |        |      |     |      |      |       |                      | 10.20.246.21[0/31]|
|            |        |      |     |      |      |       |                      | 10.20.246.22[0/31]|
|            |        |      |     |      |      |       |                      | 10.20.246.25[0/31]|
|            |        |      |     |      |      |       |                      | 10.20.246.26[0/31]|
+------------+--------+------+-----+------+------+-------+----------------------+--------------------+
|    ten1    | private| 11-20|     |      | 10   | false | 10.20.246.15[0/1-10] |                    |
|            |        |      |     |      |      |       | 10.20.246.16[0/1-10] |                    |
+------------+--------+------+-----+------+------+-------+----------------------+--------------------+
|    ten2    | private| 21-30|     |      | 10   | false | 10.20.246.15[0/11-20]|                    |
|            |        |      |     |      |      |       | 10.20.246.16[0/11-20]|                    |
+------------+--------+------+-----+------+------+-------+----------------------+--------------------+


(efa:root)root@node 2:~# efa tenant po show
+-------+-------+----+------+-----+-----------+---------+--------+-------------------+-----------
+-------------+------------+
|  Name | Tenant| ID |Speed | MTU |Negotiation |Min Link |  Lacp  |
Ports      |    State  |  Dev State  |  App State |
|       |       |    |      |     |           |  Count  | Timeout|
|       |       |    |      |     |
+-------+-------+----+------+-----+-----------+---------+--------+-------------------+-----------
```

```
+-------------+------------+
|ten1po1| ten1   |  2 |10Gbps|     |   active   |   1     |   long |
10.20.246.15[0/1] | po-created| provisioned | cfg-in-sync|
|        |       |    |      |     |            |         |        | 10.20.246.16[0/1]
|          |            |           |
+-------+-------+----+------+-----+------------+---------+--------+------------------+-----------
+-------------+------------+
|ten2po1| ten2   |  3 |10Gbps|     |   active   |   1     |   long
| 10.20.246.15[0/11]| po-created| provisioned | cfg-in-sync|
|        |       |    |      |     |            |         |        | 10.20.246.16[0/11]|
|             |            |
+-------+-------+----+------+-----+------------+---------+--------+------------------+-----------
+-------------+------------+
```

Example:

10.20.246.15	10.20.246.16
efa tenant service mirror session create --name ten1mirrorsession1 --tenant ten1 --source 10.20.246.15,po,ten1po1 **--type 10.20.246.15,po,ten1po1:port-based** --destination 10.20.246.15,po,ten1po1:10.20.246.15,eth,0/31 **--destination-type 10.20.246.15,po,ten1po1:span** --direction 10.20.246.15,po,ten1po1:both	efa tenant service mirror session create –name ten1mirrorsession2 --tenant ten1 --source 10.20.246.16,po,ten1po1 **--type 10.20.246.16,po,ten1po1:port-based** --destination 10.20.246.16,po,ten1po1:10.20.246.16,eth,0/31 **--destination-type 10.20.246.16,po,ten1po1:span** --direction 10.20.246.16,po,ten1po1:both
10.20.246.15	**10.20.246.16**
efa tenant service mirror session create –name ten2mirrorsession1 --tenant ten2 --source 10.20.246.15,po,ten2po1 **--type 10.20.246.15,po,ten2po1:port-based** --destination 10.20.246.15,po,ten2po1:10.20.246.15,eth,0/31 **--destination-type 10.20.246.15,po,ten2po1:span** --direction 10.20.246.15,po,ten2po1:both	efa tenant service mirror session create –name ten2mirrorsession2 --tenant ten2 --source 10.20.246.16,po,ten2po1 **--type 10.20.246.16,po,ten2po1:port-based** --destination 10.20.246.16,po,ten2po1:10.20.246.16,eth,0/31 **--destination-type 10.20.246.16,po,ten2po1:spa**n --direction 10.20.246.16,po,ten2po1:both

2. Verify the switch configuration on the SLX device.

10.20.246.15	10.20.246.16
SLX# show running-config monitor session monitor session 1 source port-channel 2 destination ethernet 0/31 direction both !monitor session 2 source port-channel 3 destination ethernet 0/31 direction both ! SLX# show monitor session 1 Session : 1 Type : SPAN Description : [None] State : Enabled Source Interface : Po 2 (Down) Destination Interface : Eth 0/31 (Down) Direction : Both Type : port-based SLX# show monitor session 2 Session : 2 Type : SPAN Description : [None] State : Enabled Source Interface : Po 3 (Down) Destination Interface : Eth 0/31 (Down) Direction : Both Type : port-based	SLX# show running-config monitor session monitor session 1 source port-channel 2 destination ethernet 0/31 direction both !monitor session 2 source port-channel 3 destination ethernet 0/31 direction both ! SLX# show monitor session 1 Session : 1 Type : SPAN Description : [None] State : Enabled Source Interface : Po 2 (Down) Destination Interface : Eth 0/31 (Down) Direction : Both Type : port-based SLX# show monitor session 2 Session : 2 Type : SPAN Description : [None] State : Enabled Source Interface : Po 3 (Down) Destination Interface : Eth 0/31 (Down) Direction : Both Type : port-based

Configure Flow-Based Mirroring in a Multi-Tenant Architecture

You can configure flow-based mirroring in a multi-tenant architecture.

About This Task

Follow this procedure to configure flow-based mirroring.

Procedure

1. Run the following commands to configure access control list applications on Ethernet or port channel and VLAN or virtual Ethernet:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>
   --switchport --switchport-mode trunk -ctag-range <ctag-range>
   --port <mirror-source-port-list> --po <mirror-source-po-list>

   --pp-mac-acl-in <acl-name> --pp-mac-acl-out <acl-name>
   --pp-acl-in <acl-name> --pp-ip-acl-out <acl-name>

   --np-mac-acl-in <ctag:acl-name> --np-mac-acl-out <ctag:acl-name>
   --np-ip-acl-in <ctag:acl-name> --np-ip-acl-out <ctag:acl-name>
```

2. Run the following commands to configure a mirror session:

```
efa tenant service mirror session create –name? <session-name> --tenant <tenant-name>
   --source {<device-ip>,<eth | po | vlan>,<if-name>}
   --type {<source-device-ip>,<eth | po | vlan>,<source-if-name>:<port-based | flow-
based>}

   --destination {<source-device-ip>,<eth | po | vlan>,<source-if-name> :
           <destination-device-ip>,<eth | po | vlan>,<destination-if-name}
   --destination-type {<source-device-ip>,< eth | po | vlan>,<source-if-name>:<span>}

   --direction {<source-device-ip>,< eth | po | vlan>,<source-if-name> : <tx | rx |
both>}

(efa:root)root@node-2:~# efa tenant show
+--------------+---------+------+------+------+------+-------+----------------------
+--------------------+
|     Name     |  Type   | VLAN | L2VNI| L3VNI| VRF  | Enable|
Ports        | Mirroring Ports   |
|              |         | Range| Range| Range| Count| BD    |
|                    |
+--------------+---------+------+------+------+------+-------+----------------------
+--------------------+
| sharedTenant | shared  |      |      |      |   0  | false |
| 10.20.246.15[0/31]|
|              |         |      |      |      |      |       |
| 10.20.246.16[0/31]|
|              |         |      |      |      |      |       |
| 10.20.246.21[0/31]|
|              |         |      |      |      |      |       |
| 10.20.246.22[0/31]|
|              |         |      |      |      |      |       |
| 10.20.246.25[0/31]|
|              |         |      |      |      |      |       |
| 10.20.246.26[0/31]|
+--------------+---------+------+------+------+------+-------+----------------------
+--------------------+
|     ten1     | private |11-20 |      |      |   10 | false | 10.20.246.15[0/1-10]
|                    |
|              |         |      |      |      |      |       | 10.20.246.16[0/1-10]
|                    |
+--------------+---------+------+------+------+------+-------+----------------------
+--------------------+
|     ten2     | private |21-30 |      |      |   10 | false |
10.20.246.15[0/11-20]|                    |
|              |         |      |      |      |      |       |
10.20.246.16[0/11-20]|                    |
+--------------+---------+------+------+------+------+-------+----------------------
+--------------------+


(efa:root)root@node 2:~# efa tenant po show
+---------+------+----+------+-----+------------+---------+-------+-------------------
+------------+------------+------------+
|   Name  |Tenant| ID |Speed | MTU |Negotiation |Min Link | Lacp
|        Ports      |    State   | Dev State  | App State  |
|         |      |    |      |     |            |  Count  |Timeout|
|            |            |            |
+---------+------+----+------+-----+------------+---------+-------+-------------------
+------------+------------+------------+
| ten1po1 |ten1  |  2 |10Gbps|     |   active   |   1     |  long
| 10.20.246.15[0/1] | po-created |provisioned |cfg-in-sync |
|           |      |    |      |     |            |         |       | 10.20.246.16[0/1]
|           |            |            |
+---------+------+----+------+-----+------------+---------+-------+-------------------
```

```
+------------+------------+------------+
| ten2po1 |ten2  |  3 |10Gbps|     |   active   |   1     |  long
| 10.20.246.15[0/11]| po-created |provisioned |cfg-in-sync |
|         |      |    |      |     |            |         |       |
10.20.246.16[0/11]|            |            |           |
+---------+------+----+------+-----+------------+---------+-------+------------------
+------------+------------+------------+
```

Example

```
efa tenant epg create –name
ten1epg1 –tenant ten1
  --switchport-mode trunk --po
ten1po1 --ctag-range 11
  --pp-ip-acl-in ext-ip-permit-any-
mirror-acl
  --pp-ip-acl-out  ext-ip-permit-
any-mirror-acl

efa tenant service mirror session
create –name ten1mirrorsession1 --
tenant ten1
  --source 10.20.246.15,po,ten1po1
  --type
10.20.246.15,po,ten1po1:flow-based

  --destination
10.20.246.15,po,ten1po1:10.20.246.1
5,eth,0/31
  --destination-type
10.20.246.15,po,ten1po1:span

  --direction
10.20.246.15,po,ten1po1:both

efa tenant service mirror session
create –name ten2mirrorsession1 --
tenant ten2
  --source 10.20.246.15,po,ten2po1
  --type
10.20.246.15,po,ten2po1:flow-based

  --destination
10.20.246.15,po,ten2po1:10.20.246.1
5,eth,0/31
  --destination-type
10.20.246.15,po,ten2po1:span

  --direction
10.20.246.15,po,ten2po1:both
```

```
efa tenant epg create –name
ten2epg1 –tenant ten2
  --switchport-mode trunk --po
ten2po1 --ctag-range 21
  --pp-ip-acl-in ext-ip-permit-any-
mirror-acl
  --pp-ip-acl-out ext-ip-permit-
any-mirror-acl


efa tenant service mirror session
create –name ten1mirrorsession2 --
tenant ten1
  --source 10.20.246.16,po,ten1po1
  --type
10.20.246.16,po,ten1po1:flow-based

  --destination
10.20.246.16,po,ten1po1:10.20.246.1
6,eth,0/31
  --destination-type
10.20.246.16,po,ten1po1:span

  --direction
10.20.246.16,po,ten1po1:both

efa tenant service mirror session
create –name ten2mirrorsession2 --
tenant ten2
  --source 10.20.246.16,po,ten2po1
  --type
10.20.246.16,po,ten2po1:flow-
based

  --destination
10.20.246.16,po,ten2po1:10.20.246.1
6,eth,0/31
  --destination-type
10.20.246.16,po,ten2po1:span

  --direction
10.20.246.16,po,ten2po1:both
```

3. Verify the switch configuration on the SLX device.

10.20.246.15

```
SLX# show running-config ip access-list
ip access-list extended ext-ip-permit-any-mirror-acl
 seq 10 permit ip any any mirror
!
SLX# show running-config interface Port-channel 2,3
interface Port-channel 2
 description EFA Port-channel ten1po1
 cluster-client auto
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 11
 no switchport trunk tag native-vlan
 ip access-group ext-ip-permit-any-mirror-acl in
 ip access-group ext-ip-permit-any-mirror-acl out
 no shutdown
!
interface Port-channel 3
 description EFA Port-channel ten2po1
 cluster-client auto
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 21
 no switchport trunk tag native-vlan
 ip access-group ext-ip-permit-any-mirror-acl in
 ip access-group ext-ip-permit-any-mirror-acl out
 no shutdown
!
SLX#
```

10.20.246.16

```
SLX# show running-config ip access-list
ip access-list extended ext-ip-permit-any-mirror-acl
 seq 10 permit ip any any mirror
!
SLX# show running-config interface Port-channel 2,3
interface Port-channel 2
 description EFA Port-channel ten1po1
 cluster-client auto
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 11
 no switchport trunk tag native-vlan
 ip access-group ext-ip-permit-any-mirror-acl in
 ip access-group ext-ip-permit-any-mirror-acl out
 no shutdown
!
interface Port-channel 3
 description EFA Port-channel ten2po1
 cluster-client auto
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 21
 no switchport trunk tag native-vlan
 ip access-group ext-ip-permit-any-mirror-acl in
 ip access-group ext-ip-permit-any-mirror-acl out
 no shutdown
!
SLX#
```

10.20.246.15

```
SLX# show running-config monitor session
monitor session 1
 source port-channel 2 destination ethernet 0/31 direction both
!monitor session 2
 source port-channel 3 destination ethernet 0/31 direction both
!
SLX# show monitor session 1
Session                 : 1
Type                    : SPAN
Description             : [None]
State                   : Enabled
Source Interface        : Po 2 (Down)
Destination Interface   : Eth 0/31 (Down)
```

10.20.246.16

```
SLX# show running-config monitor session
monitor session 1
 source port-channel 2 destination ethernet 0/31 direction both
!monitor session 2
 source port-channel 3 destination ethernet 0/31 direction both
!
SLX# show monitor session 1
Session                 : 1
Type                    : SPAN
Description             : [None]
State                   : Enabled
Source Interface        : Po 2 (Down)
Destination Interface   : Eth 0/31 (Down)
```

```
Direction                : Both
Type                     : flow-based

SLX# show monitor session 2
Session                  : 2
Type                     : SPAN
Description              : [None]
State                    : Enabled
Source Interface         : Po 3
(Down)
Destination Interface    : Eth 0/31
(Down)
Direction                : Both
Type                     : flow-based
```

```
Direction                : Both
Type                     : flow-based

SLX# show monitor session 2
Session                  : 2
Type                     : SPAN
Description              : [None]
State                    : Enabled
Source Interface         : Po 3
(Down)
Destination Interface    : Eth 0/31
(Down)
Direction                : Both
Type                     : flow-based
```

Access Control List and Data Consistency Support

XCO provisions the Access Control List (ACL) internally because there is no ACL CRUD support available from the XCO side.

Full-fledged data consistency support (any drift in the ACL rules is identified and reconciled) is available when the ACL CRUD is supported via XCO.

Configure VLAN-Based Mirroring in a Multi-Tenant Architecture

You can configure VLAN-based mirroring in a multi-tenant architecture.

Before You Begin

VLAN-based mirroring applies only to VLAN-based tenants and not to BD (bridge domain)-based tenants.

About This Task

Follow this procedure to configure VLAN-based mirroring.

Procedure

1. Run the following commands to configure access control list applications on Ethernet or Port channel and VLAN or Virtual Ethernet:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>

   --switchport --switchport-mode trunk -ctag-range <ctag-range>
   --port <mirror-source-port-list> --po <mirror-source-po-list>

   --pp-mac-acl-in <acl-name> --pp-mac-acl-out <acl-name>
   --pp-ip-acl-in <acl-name> --pp-ip-acl-out <acl-name>

   --np-mac-acl-in <ctag:acl-name> --np-mac-acl-out <ctag:acl-name>
   --np-ip-acl-in <ctag:acl-name> --np-ip-acl-out <ctag:acl-name>
```

2. Run the following commands to configure a mirror session:

```
efa tenant service mirror session create -name <session-name> --tenant <tenant-name>

   --source {<device-ip>,<eth | po | vlan>,<if-name>}
   --type {<source-device-ip>,<eth | po | vlan>,<source-if-name>:<port-based | flow-
based>}

   --destination {<source-device-ip>,<eth | po | vlan>,<source-if-name> :
               <destination-device-ip>,<eth | po | vlan>,<destination-if-name}
```

```
   --destination-type {<source-device-ip>,< eth | po | vlan>,<source-if-name>:<span>}

   --direction {<source-device-ip>,< eth | po | vlan>,<source-if-name> : <tx | rx |
both>}


(efa:root)root@node-2:~# efa tenant show
+-------+-------+------+------+------+------+------+---------------------
+------------------+
|Name   | Type  | VLAN | L2VNI| L3VNI| VRF
|Enable|        Ports        |  Mirroring Ports |
|       |       | Range| Range| Range| Count|BD
|                      |                  |
+-------+-------+------+------+------+------+------+---------------------
+------------------+
|shared |Shared |      |      |      |   0
|false |                     |10.20.246.16[0/31]|
|Tenant |       |      |      |      |      |
|                     |10.20.246.21[0/31]|
|       |       |      |      |      |      |
|                     |10.20.246.22[0/31]|
|       |       |      |      |      |      |
|                     |10.20.246.25[0/31]|
|       |       |      |      |      |      |
|                     |10.20.246.26[0/31]|
+-------+-------+------+------+------+------+------+---------------------
+------------------+
| ten1  |private| 11-20|      |      |   10
|false |10.20.246.15[0/1-10] |                  |
|       |       |      |      |      |      |
|10.20.246.16[0/1-10] |                  |
|       |       |      |      |      |      |
|10.20.246.21[0/1-10] |                  |
|       |       |      |      |      |      |
|10.20.246.22[0/1-10] |                  |
+-------+-------+------+------+------+------+------+---------------------
+------------------+
| ten2  |private| 21-30|      |      |   10
|false |10.20.246.15[0/11-20]|                 |
|       |       |      |      |      |      |
|10.20.246.16[0/11-20]|                  |
|       |       |      |      |      |      |
|10.20.246.21[0/11-20]|                  |
|       |       |      |      |      |      |
|10.20.246.22[0/11-20]|                  |
+-------+-------+------+------+------+------+------+---------------------
+------------------+


(efa:root)root@node 2:~# efa tenant po show
+--------+-------+-----------+-----+-----------+----------+-------+-------------------
+-----------+-------------+-------------+
|  Name  |Tenant |ID | Speed | MTU |Negotiation|
Min Link | Lacp  |        Ports      |    State  |  Dev State  |  App State  |
|        |       |   |       |     |
|   Count  |Timeout|                   |           |             |             |
+--------+-------+-----------+-----+-----------+----------+-------+-------------------
+-----------+-------------+-------------+
|ten1po1 |ten1   | 2 | 10Gbps|     |   active
|    1     |  long | 10.20.246.15[0/1] | po-created| provisioned | cfg-in-sync |
|        |       |   |       |     |
|          |       | 10.20.246.16[0/1] |           |             |             |
+--------+-------+-----------+-----+-----------+----------+-------+-------------------
+-----------+-------------+-------------+
```

```
|ten2po1 |ten2   | 3 | 10Gbps|     |   active
|    1     |  long | 10.20.246.15[0/11]| po-created| provisioned | cfg-in-sync |
|        |       |   |       |     |
|          |       | 10.20.246.16[0/11]|           |             |             |
+--------+-------+-----------+-----+-----------+----------+-------+------------------
+-----------+-------------+-------------+
|ten1po2 |ten1   | 2 | 10Gbps|     |   active
|    1     |  long | 10.20.246.21[0/1] | po-created| provisioned | cfg-in-sync |
|        |       |   |       |     |
|          |       | 10.20.246.22[0/1] |           |             |             |
+--------+-------+-----------+-----+-----------+----------+-------+------------------
+-----------+-------------+-------------+
|ten2po2 |ten2   | 3 | 10Gbps|     |   active
|    1     |  long | 10.20.246.21[0/11]| po-created| provisioned | cfg-in-sync |
|        |       |   |       |     |
|          |       | 10.20.246.22[0/11]|           |             |             |
+--------+-------+---+-------+-----+-----------+----------+-------+------------------
+-----------+-------------+-------------+
```

Example

```
efa tenant epg create –name
ten1epg1 –tenant ten1
   --switchport-mode trunk --po
ten1po1,ten1po2  --ctag-range 11
   --np-mac-acl-in 11:ext-mac-
permit-any-mirror-acl
   --np-mac-acl-out 11:ext-mac-
permit-any-mirror-acl


efa tenant service mirror session
create –name ten1mirrorsession1 --
tenant ten1
   --source vlan,11
   --type vlan,11:flow-based
   --destination-type vlan,11:span
   --destination
vlan,11:10.20.246.15,eth,0/31
   --direction vlan,11:both
```

```
efa tenant epg create –name
ten2epg1 –tenant ten2
   --switchport-mode trunk --po
ten2po1,ten2po2 --ctag-range 21
   --np-mac-acl-in 21:ext-mac-
permit-any-mirror-acl
   --np-mac-acl-out 21:ext-mac-
permit-any-mirror-acl


efa tenant service mirror session
create –name ten2mirrorsession1 --
tenant ten2
   --source vlan,21
   --type vlan,21:flow-based
   --destination-type vlan,21:span
   --destination
vlan,21:10.20.246.16,eth,0/31
   --direction vlan,21:both
```

3. Verify the switch configuration on the SLX device.

10.20.246.15

```
SLX# show running-config mac access-list
mac access-list extended ext-mac-permit-any-mirror-acl
 seq 10 permit any any mirror
!
SLX#

SLX# show running-config vlan 11,21
vlan 11
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
vlan 21
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
SLX#
```

10.20.246.16

```
SLX# show running-config mac access-list
mac access-list extended ext-mac-permit-any-mirror-acl
 seq 10 permit any any mirror
!
SLX#

SLX# show running-config vlan 11,21
vlan 11
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
vlan 21
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
SLX#
```

10.20.246.21

```
SLX# show running-config mac access-list
mac access-list extended ext-mac-permit-any-mirror-acl
 seq 10 permit any any mirror
!
SLX#

SLX# show running-config vlan 11,21
vlan 11
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
vlan 21
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
SLX#
```

10.20.246.22

```
SLX# show running-config mac access-list
mac access-list extended ext-mac-permit-any-mirror-acl
 seq 10 permit any any mirror
!
SLX#

SLX# show running-config vlan 11,21
vlan 11
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
vlan 21
 description Tenant L2 Extended VLAN
 mac access-group ext-mac-permit-any-mirror-acl in
 mac access-group ext-mac-permit-any-mirror-acl out
!
SLX#
```

10.20.246.15-16

```
SLX# show running-config monitor session
monitor session 1
 source vlan 11 destination ethernet 0/31 direction both flow-based
!
monitor session 2
 source vlan 21 destination ethernet 0/31 direction both flow-based
!SLX#
```

10.20.246.21-22

```
SLX# show running-config monitor session
monitor session 1
 source vlan 11 destination ethernet 0/31 direction both flow-based
!
monitor session 2
 source vlan 21 destination ethernet 0/31 direction both flow-based
!SLX
```

```
SLX# show monitor session 1
Session                  : 1
Type                     : SPAN
Description              : [None]
State                    : Enabled
Source Interface         : Vlan 11
Destination Interface    : Eth 0/31
(Down)
Direction                : Both
Type                     : flow-based

SLX# show monitor session 2
Session                  : 2
Type                     : SPAN
Description              : [None]
State                    : Enabled
Source Interface         : Vlan 21
Destination Interface    : Eth 0/31
(Down)
Direction                : Both
Type                     : flow-based
SLX#
```

```
#SLX# show monitor session 1
Session                  : 1
Type                     : SPAN
Description              : [None]
State                    : Enabled
Source Interface         : Vlan 11
Destination Interface    : Eth 0/31
(Down)
Direction                : Both
Type                     : flow-based

SLX# show monitor session 2
Session                  : 2
Type                     : SPAN
Description              : [None]
State                    : Enabled
Source Interface         : Vlan 21
Destination Interface    : Eth 0/31
(Down)
Direction                : Both
Type                     : flow-based
SLX#
```

Configure ICL Port Mirroring in a Multi-Tenant Architecture

You can configure an ICL port mirroring in a multi-tenant architecture.

About This Task

Follow this procedure to configure an ICL port mirroring.

Procedure

1. Run the following commands to configure access control list applications on Ethernet or Port channel and VLAN or Virtual Ethernet:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>

    --type port-profile
    --po <mirror-source-po-list>

    --pp-ipv6-acl-in <acl-name>
    --pp-ip-acl-in <acl-name> --pp-ip-acl-out <acl-name>
```

2. Run the following commands to configure a mirror session:

```
efa tenant service mirror session create -name <session-name> --tenant <tenant-name>
    --source {<device-ip>,<eth | po | vlan>,<if-name>}
    --type {<source-device-ip>,<eth | po | vlan>,<source-if-name>:<port-based | flow-
based>}

    --destination-type {<source-device-ip>,< eth | po | vlan>,<source-if-name>:<span>}
    --destination {<source-device-ip>,<eth | po | vlan>,<source-if-name> :
            <destination-device-ip>,<eth | po | vlan>,<destination-if-name}
    --direction {<source-device-ip>,< eth | po | vlan>,<source-if-name> : <tx | rx |
both>}


(efa:root)root@node-2:~# efa tenant show
+-------+------+------+------+------+-------+----------------------------
+-------------------+
|Name   | Type | VLAN | L2VNI|L3VNI | VRF  |Enable |
Ports            |  Mirroring Ports  |
|       |      | Range| Range|Range | Count|BD     |
```

```
|                   |
+-------+------+------+------+------+------+-------+----------------------------
+--------------------+
|shared |Shared|      |      |      |   0  |false  |10.20.246.15[0/46-47]
|10.20.246.15[0/31] |
|Tenant |      |      |      |      |      |       |10.20.246.16[0/46-47]
|10.20.246.16[0/31] |
|       |      |      |      |      |      |       |10.20.246.21[0/9-10,0/46-48]|
10.20.246.21[0/31] |
|       |      |      |      |      |      |       |10.20.246.22[0/9-10,0/46-48]|
10.20.246.22[0/31] |
|       |      |      |      |      |      |       |
|10.20.246.25[0/31] |
|       |      |      |      |      |      |       |
|10.20.246.26[0/31] |
+-------+------+------+------+------+------+-------+----------------------------
+--------------------+


(efa:root)root@node 2:~# efa tenant po show
+-------+------+--+------+---+-----------+--------+-------+----------------+-----------
+------------+-----------+
|  Name |Tenant|ID|Speed |MTU|Negotiation|Min Link| Lacp  |
Ports      |   State   | Dev State  | App State |
|       |      |  |      |   |           |Count   |Timeout|
|          |           |           |
+-------+------+--+------+---+-----------+--------+-------+----------------+-----------
+------------+-----------+
|ten1po1|ten1  |64|10Gbps|   |   active  |   1    |  long |
10.20.246.15    |po-created |provisioned |cfg-in-sync|
|       |      |  |      |   |           |        |       |
[0/46-47]       |           |            |           |
|       |      |  |      |   |           |        |       |
10.20.246.16    |           |            |           |
|       |      |  |      |   |           |        |       |
[0/46-47]       |           |            |           |
+-------+------+--+------+---+-----------+--------+-------+----------------+-----------
+------------+-----------+
|ten2po1|ten2  |64|10Gbps|   |   active  |   1    |  long |
10.20.246.21    |po-created |provisioned |cfg-in-sync|
|       |      |  |      |   |           |        |       |
0/9-10,0/46-48] |           |            |           |
|       |      |  |      |   |           |        |       |
10.20.246.22|   |           |            |           |
|       |      |  |      |   |           |        |       |
[0/9-10,0/46-48]|           |            |           |
+-------+------+--+------+---+-----------+--------+-------+----------------+-----------
+------------+-----------+
```

Example

efa tenant epg create –name ten1epg1 –tenant ten1 --type port-profile
  --po ten1po1
  --pp-ipv6-acl-in ext-ipv6-permit-any-mirror-acl
	efa tenant epg create –name ten1epg2 –tenant ten1 --type port-profile
	  --po ten1po2
	  --pp-ipv6-acl-in ext-ipv6-permit-any-mirror-acl

efa tenant service mirror session create –name mirrorsession1 --tenant ten1
  --source 10.20.246.15,po,ten1po1
  --type 10.20.246.15,po,ten1po1:port-based
  --destination 10.20.246.15,po,ten1po1:10.20.246.15,eth,0/31
  --destination-type 10.20.246.15,po,ten1po1:span
  --direction 10.20.246.15,po,ten1po1:**tx**
	efa tenant service mirror session create –name mirrorsession3 --tenant ten1
	  --source 10.20.246.21,po,ten1po2
	  --type 10.20.246.21,po,ten1po2:port-based
	  --destination 10.20.246.21,po,ten1po2:10.20.246.21,eth,0/31
	  --destination-type 10.20.246.21,po,ten1po2:span
	  --direction 10.20.246.21,po,ten1po2:**tx**

efa tenant service mirror session create –name mirrorsession2 --tenant ten1
  --source 10.20.246.15,po,ten1po1
  --type 10.20.246.15,po,ten1po1:flow-based
  --destination 10.20.246.15,po,ten1po1:10.20.246.15,eth,0/31
  --destination-type 10.20.246.15,po,ten1po1:span
  --direction 10.20.246.15,po,ten1po1:**rx**
	efa tenant service mirror session create –name mirrorsession4 --tenant ten1
	  --source 10.20.246.21,po,ten1po2
	  --type 10.20.246.21,po,ten1po2:flow-based
	  --destination 10.20.246.21,po,ten1po2:10.20.246.21,eth,0/31
	  --destination-type 10.20.246.21,po,ten1po2:span
	  --direction 10.20.246.21,po,ten1po2:**rx**

3. Verify the switch configuration on the SLX device.

10.20.246.15

```
SLX# show running-config ipv6
access-list
ipv6 access-list extended ext-ipv6-
permit-any-mirror-acl
 seq 10 permit ipv6 any any mirror
!
SLX#

SLX# show running-config int po 64
interface Port-channel 64
 mtu 9216
 description MCTPeerInterface
 ip address 10.20.20.3/31
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 no shutdown
!
SLX#
```

10.20.246.16

```
SLX# show running-config ipv6
access-list
ipv6 access-list extended ext-ipv6-
permit-any-mirror-acl
 seq 10 permit ipv6 any any mirror
!
SLX#

SLX# show running-config int po 64
interface Port-channel 64
 mtu 9216
 description MCTPeerInterface
 ip address 10.20.20.2/31
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 no shutdown
!
SLX#
```

10.20.246.21

```
SLX# show running-config ipv6
access-list
ipv6 access-list extended ext-ipv6-
permit-any-mirror-acl
 seq 10 permit ipv6 any any mirror
!
SLX#

SLX# show running-config int po 64
interface Port-channel 64
 mtu 9216
 description MCTPeerInterface
 ip address 10.20.20.3/31
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 no shutdown
!
SLX#
```

10.20.246.22

```
SLX# show running-config ipv6
access-list
ipv6 access-list extended ext-ipv6-
permit-any-mirror-acl
 seq 10 permit ipv6 any any mirror
!
SLX#

SLX# show running-config int po 64
interface Port-channel 64
 mtu 9216
 description MCTPeerInterface
 ip address 10.20.20.2/31
 ipv6 access-group ext-ipv6-permit-
any-mirror-acl in
 no shutdown
!
SLX#
```

10.20.246.15

```
SLX# show running-config monitor
session
monitor session 1
 source port-channel 64
destination ethernet 0/31
direction tx
!
monitor session 2
 source port-channel 64
destination ethernet 0/31
direction rx flow-based
!
SLX# show monitor session 1
Session                : 1
Type                   : SPAN
Description            : [None]
State                  : Enabled
Source Interface       : Po 64 (Up)
Destination Interface  : Eth 0/31
(Down)
```

10.20.246.21

```
SLX# show running-config monitor
session
monitor session 1
 source port-channel 64
destination ethernet 0/31
direction tx
!
monitor session 2
 source port-channel 64
destination ethernet 0/31
direction rx flow-based
!
SLX# show monitor session 1
Session                : 1
Type                   : SPAN
Description            : [None]
State                  : Enabled
Source Interface       : Po 64 (Up)
Destination Interface  : Eth 0/31
(Down)
```

Direction : Tx **Type : port-based** SLX# show monitor session 2 Session : 2 Type : SPAN Description : [None] State : Enabled Source Interface : Po 64 (Up) Destination Interface : **Eth 0/31 (Down)** Direction : Rx Type : **flow-based**	Direction : Tx **Type : port-based** SLX# show monitor session 2 Session : 2 Type : SPAN Description : [None] State : Enabled Source Interface : Po 64 (Up) **Destination Interface : Eth 0/31 (Down)** Direction : Rx **Type : flow-based**

Configure Fabric Non-ICL Ports as Mirror Source

You can configure fabric non-ICL port as mirror source.

About This Task

Follow this procedure to configure fabric non-ICL port mirror source.

Mirror the traffic from the spine and super spine ports onto the mirror destination port. The provisioning model is inline with the ICL port channel mirroring.

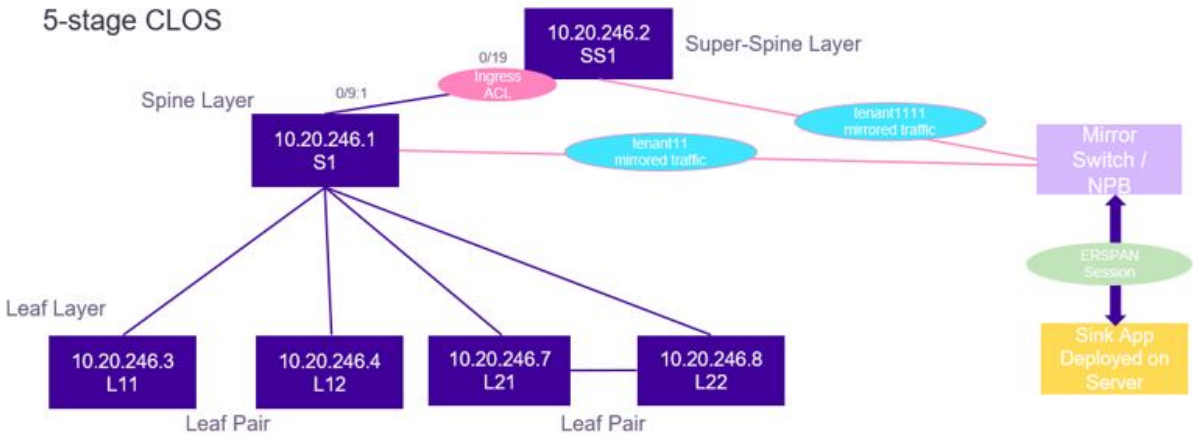

Figure 31: 5-stage Clos topology

Note

1. Spine and super spine ports can be a member of the shared tenant only and not the private tenant.
2. Spine and super spine ports can be a member of the port profile EPG only and not any other EPG.
3. You cannot create a port channel using the spine and super spine ports.
4. You cannot apply any other configurations on the spine or super spine.

Procedure

1. Create a shared tenant using the spine and super spine ports.

2. Create an EPG port profile with spine and super spine ports as endpoints of an EPG. Ensure that the port profile EPG is under the shared tenant. This creates an ACL application on the spine and super spine ports for flow-based mirroring.

```
efa tenant epg create --name <epg-name> --tenant <tenant-name>
     --type port-profile --port <spine-and-super-spine-mirror-ports>
     --pp-ip-acl-in <acl-name> --pp-ip-acl-out <acl-name> --pp-ipv6-acl-in <acl-name>
```

3. Create a mirror session using spine and super spine ports as a mirror source.

```
efa tenant service mirror session create -name <session-name> --tenant <tenant-name>
    --source {<device-ip>,<eth | po | vlan>,<if-name>}
    --type {<source-device-ip>,<eth | po | vlan>,<source-if-name>:<port-based | flow-
based>}
    --destination-type {<source-device-ip>,< eth | po | vlan>,<source-if-name>:<span>}
    --destination {<source-device-ip>,<eth | po | vlan>,<source-if-name> :
                  <destination-device-ip>,<eth | po | vlan>,<destination-if-name>
    --direction {<source-device-ip>,< eth | po | vlan>,<source-if-name> : <tx | rx |
both>}
```

For example,

```
(efa:root)root@node-2:~# efa tenant show
+---------+------+-------+-----+------+-----+------+-------------------------
+--------------------+
|  Name   | Type |VLAN   |L2VNI| L3VNI|VRF  |Enable|
Ports           |Mirror Destination  |
|         |      |Range  |Range| Range|Count|BD    |
|Ports               |
+---------+------+-------+-----+------+-----+------+-------------------------
+--------------------+
|tenant11 |shared|100-103|     |      |10   |false |10.20.246.1
|10.20.246.1[0/17-18]|
|         |      |       |     |      |     |      |
[0/10-11,0/31-32,0/9:1-4]|10.20.246.1[0/17-18]|
|         |      |       |     |      |     |      |
10.20.246.3[0/10-15]     |10.20.246.3[0/1-9]  |
|         |      |       |     |      |     |      |
10.20.246.4[0/1-15]      |                    |
+---------+------+-------+-----+------+-----+------+-------------------------
+--------------------+
|tenant111|shared|104-105|     |      |10   |false |10.20.246.1
|[0/13-15,0/19,      |
|         |      |       |     |      |     |      |
[0/14-15,0/13:1-4]       |0/12:1-4]           |
|         |      |       |     |      |     |      |
10.20.246.2              |                    |
+---------+------+-------+-----+------+-----+------+-------------------------
+--------------------+
(efa:root)root@node-2:~# efa tenant epg show
+-------+---------+-----------+-----------+--+---------+-----------+-----+---+------
+--------------------+
| Name  | Tenant  |    Type    |   Ports
|PO|Switch   |Native Vlan|Ctag |Vrf| L3Vni|       State       |
|       |         |           |           |  |POrt Mode|
Tagging  |Range|   |      |                  |
+-------+---------+-----------+-----------+--+---------+-----------+-----+---+------
+--------------------+
|epgv421|tenant111|port-profile|10.20.246.2|  |         |   false
|     |   |      |epg-with-port-group|
|       |         |           |[0/19]     |  |
|           |     |   |      |                   |
+-------+---------+-----------+-----------+--+---------+-----------+-----+---+------
+--------------------+

efa tenant service mirror session create --name
```

```
m2 --tenant "tenant11" --source 10.20.246.1,eth,0/9:1
       --type 10.20.246.1,eth,0/9:1:port-based
       --destination 10.20.246.1,eth,0/9:1:10.20.246.1,eth,0/17
       --destination-type 10.20.246.1,eth,0/9:1:span
       --direction 10.20.246.1,eth,0/9:1:rx
efa tenant service mirror session create --name
m3 --tenant tenant11 --source 10.20.246.1,eth,0/9:1
       --type 10.20.246.1,eth,0/9:1:port-based
       --destination 10.20.246.1,eth,0/9:1:10.20.246.1,eth,0/17
       --destination-type 10.20.246.1,eth,0/9:1:span
       --direction 10.20.246.1,eth,0/9:1:tx
efa tenant epg create –name epgv421 –tenant tenant111 --type port-profile
       --port 10.20.246.2[0/19] --pp-ipv6-acl-in ext-ipv6-permit-any-mirror-acl
efa tenant service mirror session create --name
ms3 --tenant tenant111 --source 10.20.246.2,eth,0/19
       --type 10.20.246.2,eth,0/19:flow-based
       --destination 10.20.246.2,eth,0/19:10.20.246.2,eth,0/18
       --destination-type 10.20.246.2,eth,0/19:span
       --direction 10.20.246.2,eth,0/19:tx

efa tenant service mirror session create --name
ms4 --tenant tenant111 --source 10.20.246.2,eth,0/19
    --type 10.20.246.2,eth,0/19:flow-based
    --destination 10.20.246.2,eth,0/19:10.20.246.2,eth,0/18
    --destination-type 10.20.246.2,eth,0/19:span
    --direction 10.20.246.2,eth,0/19:rx
```

4. Verify the switch configuration on the SLX device.

10.20.246.1 [PORT-BASED MIRRORING]

```
SLX# show running-config monitor
session
monitor session 1
 source ethernet 0/9:1 destination
ethernet 0/17 direction rx
!
monitor session 2
 source ethernet 0/9:1 destination
ethernet 0/17 direction tx
!

SLX# show monitor session 1
Session                 : 1
Type                    : SPAN
Description             : [None]
State                   : Enabled
Source Interface        : Eth 0/9:1
Destination Interface   : Eth 0/17
(Down)
Direction               : rx
Type                    : port-based

SLX# show monitor session 2
Session                 : 2
Type                    : SPAN
Description             : [None]
State                   : Enabled
Source Interface        : Eth 0/9:1
Destination Interface   : Eth 0/17
(Down)
Direction               : tx
Type                    : port-based
```

10.20.246.2 [FLOW-BASED MIRRORING]

```
SLX# show running-config ipv6
access-list
ipv6 access-list extended ext-ipv6-
permit-any-mirror-acl
 seq 10 permit ipv6 any any mirror
!

SLX# show running-config interface
ethernet 0/19
interface ethernet 0/19
   ip address 10.10.10.1/31
   ipv6 access-group ext-ipv6-
permit-any-mirror-acl in
   no shutdown
!

SLX# show running-config monitor
session
monitor session 1
 source ethernet 0/19 destination
ethernet 0/18 direction tx flow-
based
!
monitor session 2
 source ethernet 0/19 destination
ethernet 0/18 direction rx flow-
based
!
SLX# show monitor session 1
Session                 : 1
Type                    : SPAN
Description             : [None]
State                   : Enabled
Source Interface        : Eth 0/19
Destination Interface   : Eth 0/18
(Down)
Direction               : tx
Type                    : flow-based

SLX# show monitor session 2
Session                 : 2
Type                    : SPAN
Description             : [None]
State                   : Enabled
Source Interface        : Eth 0/19
Destination Interface   : Eth 0/18
(Down)
Direction               : rx
Type                    : flow-based
```

Exclusion of VLANs and Bridge from Cluster Instance

XCO excludes VLANs and bridge domains used in the Layer 3 hand-off (toward the external gateway) endpoint group from the cluster instance by configuring `member`

`vlan remove <vlan-range>` and `member bridge-domain remove <bd-range>` under the cluster instance.

During XCO upgrade, XCO marks all the VLANs and Bridge Domains (BD) used in `l3-hand-off` EPGs with the intended `member vlan remove <vlan-range>` and `member bridge-domain remove <bd-range>` configuration and shows as configuration drift. On reconciliation of the drift, XCO pushes `member vlan remove <vlan-range>` and `member bridge-domain remove <bd-range>` configuration under the cluster.

XCO Provisioning

```
# efa tenant create --name tenant1 --port 10.24.80.134[0/1-10],10.24.80.135[0/1-10]
--vlan-range 2001-2010

# efa tenant po create --name po1 --tenant tenant1 --port
10.24.80.134[0/1],10.24.80.135[0/1]
--speed 10Gbps --negotiation active

# efa tenant epg create --name L3HandoffEPG1Ten1 --tenant tenant1 --ctag-range 2001-2003
--switchport-mode trunk --po po1 --type l3-hand-off
```

```
Device1 # show run interface Port-channel 1

interface Port-channel 1
cluster-client auto
switchport
switchport mode trunk
switchport trunk allowed vlan add 2001-2003
no switchport trunk tag native-vlan
no shutdown
!
Device1# show running config-evpn
evpn-fabric1
route-target both auto ignore-as
rd auto
duplicate-mac-timer 5 max-count 3
!
Device1# show running-config cluster
cluster fabric1-cluster-1
peer 10.20.20.5
peer-interface Port-channel 64
peer-keepalive
auto
!
member vlan-all
member vlan remove 2001-2003
member bridge-domain all
!
```

```
Device2 # show run interface Port-channel 1

interface Port-channel 1
cluster-client auto
switchport
switchport mode trunk
switchport trunk allowed vlan add 2001-2003
no switchport trunk tag native-vlan
no shutdown
!
Device2# show running config-evpn
evpn-fabric1
```

```
route-target both auto ignore-as
rd auto
duplicate-mac-timer 5 max-count 3
!
Device2# show running-config cluster
cluster fabric1-cluster-1
peer 10.20.20.5
peer-interface Port-channel 64
peer-keepalive
auto
!
member vlan-all
member vlan remove 2001-2003
member bridge-domain all
!
```

In-flight Transaction Recovery

XCO can recover in-flight (in-progress) transactions after a service restart or high-availability failover.

Overview

In-flight transactions are those that are outstanding in the execution log after a restart or a failover. After a service restart or high-availability failover, XCO recovers in-flight transactions by rolling them backward or rolling them forward.

- When transactions are rolled backward, the requested action is incomplete.
- When transactions are rolled forward, the requested action is completed.

By default, the in-flight transaction recovery feature enables the automatic recovery of Day-1 through Day-N operations for tenant-related configurations. You can use the **`efa system feature update --inflight-transaction-auto-recovery disable`**

command to disable the feature. The following table describes the recovery strategy when the feature is enabled:

Table 19: Recovery strategy

Operation type	Commands	Strategy
Create operations	• **efa tenant create** • **efa tenant epg create** • **efa tenant po create** • **efa tenant service bgp peer create** • **efa tenant service bgp peer-group create** • **efa tenant vrf create**	Roll back
Delete operations	• **efa tenant delete** • **efa tenant epg delete** • **efa tenant po delete** • **efa tenant service bgp peer delete** • **efa tenant service bgp peer-group delete** • **efa tenant vrf delete**	Roll forward
Update with add operations, such as port-add, ctag-range-add, and vrf-add	• **efa tenant update** • **efa tenant epg update** • **efa tenant po update** • **efa tenant service bgp peer update** • **efa tenant service bgp peer-group update** • **efa tenant vrf update**	Roll back
Update with delete operations, such as port-delete, vrf-delete, and ctag-range-delete	• **efa tenant update** • **efa tenant epg update** • **efa tenant po update** • **efa tenant service bgp peer update** • **efa tenant service bgp peer-group update** • **efa tenant vrf update**	Roll forward

Consider the following expected behaviors for in-flight transaction recovery:

- During operations that take a long time, such as drift and reconcile and firmware downloads, tenant operations and recovery operations are blocked.
- When multiple transactions are pending in the execution log after a restart or a failover, recovery occurs in the order in which the operations appear in the execution log.
- If a service restart or high availability failover occurs during transaction recovery, then the status of those recovery operations is changed to a normal state. For example, if a restart occurs during the rollback of an endpoint group (EPG), the

status changes to delete-pending. There is no automatic recovery of interrupted recovery transactions. You must manually verify and address the status of such operations.

Important

Day-0 and administrative operations (those for the Inventory Services and Fabric Services) are not recovered automatically. If these operations are interrupted by a service restart or a failover, you must manually redo the operations.

Examples

The following example enables automatic in-flight transaction recovery:

```
efa system feature update --inflight-transaction-auto-recovery enable
Feature Setting Updated Successful
--- Time Elapsed: 634.557118ms ---
```

The following example disables automatic in-flight transaction recovery:

```
efa system feature update --inflight-transaction-auto-recovery disable
Feature Setting Updated Successful
--- Time Elapsed: 634.557125ms ---
```

Scalability

Use this topic to learn about the scalability of tenant configuration DRC timeout and REST request timeout.

Scaled REST Request Timeout

When you run a scaled XCO tenant REST request (which takes more than 15 minutes), it fails with the following error:

```
Service is not available or internal server error has occurred, please
try again later.
```

Run the **show** command to verify the successful completion of failed REST request.

Scaled DRC Timeout

Tenant DRC Behavior	Tenant DRC Behavior in case of Scale Config
• When you run the **`efa inventory drift-reconcile execute`** command, the tenant configuration drift and reconciliation starts with other services (for example, inventory, fabric, and policy). • During the tenant DRC, the tenant drift is identified, and then the same drift is reconciled. The timeout for both tenant configuration drift identification and reconciliation is 15 minutes. • The inventory drift-reconcile execute command waits for the tenant DRC response for 15 minutes. If the response is not received within 15 minutes, then the DRC fails with the reason code `tenant-dr-timeout`. • The functionality is applicable for manual DRC and auto DRC (triggers when the device is reloaded in maintenance mode).	• In case of scaled tenant configuration drift (for example, 2000 VE or PO drifted and 100 VRFs drifted), the tenant can take more than 15 minutes for drift identification and reconciliation. Therefore, the DRC fails with the status `tenant-dr-timeout`. • The tenant configuration drift identification and reconciliation run to completion in the background even though the DRC fails with the status `tenant-dr-timeout`. • The subsequent DRC reflects the correct status of the drift.

Policy Service Provisioning

Policy Service Provisioning Overview

Policy Service in XCO manages and configures policies, such as IP prefix lists and route maps, on fabric devices.

Note
Brownfield deployment does not support the configuration managed by Policy Service.

Database, REST API and inter-service communication

Policy Service has its own database and provides REST APIs for its clients to configure and manage entities. It also registers with RabbitMQ to receive or publish messages.

Inventory Service interactions

Policy Service subscribes to the Inventory Service to receive events including device registration, device deletion, and changes to previously identified IP prefix lists and route maps.

During initialization or startup of the Policy Service, it fetches the essential entities, like device info, using REST APIs to populate its database.

Policy Service supports Drift and Reconcile (DRC) to receive and process the DRC events.

Prefix List

A prefix list allows routing systems to determine which routes to accept when they peer with their neighbors. A prefix list includes IP prefixes with a match criteria that allows or denies route redistribution. Prefix lists may contain one or more ordered entries which are processed sequentially.

XCO enables you to create, delete, and list IPv4 prefix list on a set of fabric devices. If you have not specified the ge or le value, then entry is matched with an exact prefix.

Configure IP Prefix List on Devices

Policy service supports configuration of IP prefix list for IPv4 and IPv6.

About This Task

Follow this procedure to configure IP prefix list

Procedure

1. Run the following command to configure the IPv4 prefix list:

```
efa policy prefix-list create ?
Flags:
  --type string          Type of prefix-list. Valid types is ipv4|ipv6
  --name string          Name of Prefix list
  --rule stringArray     Rule in format  seq[seq-num], action[permit/deny],
prefix[IPv4 prefix|IPv6 prefix],ge[prefix-len],le[prefix-len]. Example: seq [5],
action[permit], prefix [10.0.0.0/8|2001:db8: :/32],ge[10], le[24]
```

Note

Use the `ge` and `le` keywords to specify the range of the prefix length for exact match. Exact match is assumed when neither `ge` nor `le` is specified.

The following example creates an IPv4 prefix list:

```
efa policy prefix-list create --name prefix_v4 --type ipv4 --rule
seq[5],action[permit],prefix[10.0.0.0/8],ge[16]
```

The following example creates an IPv6 prefix list:

```
efa policy prefix-list create --type ipv6 --name prefix_1_in -rule "seq[11],
action[permit], prefix[2001:db8::/48]"

Name: prefix_1_in
+------+---------+--------+---------------+----+----+---------+
| Type | Seq num | Action |    Prefix     | Ge | Le | Status  |
+------+---------+--------+---------------+----+----+---------+
| ipv6 | 11      | permit | 2001:db8::/48 |    |    | Success |
+------+---------+--------+---------------+----+----+---------+
Prefix-list details
```

2. Run the following command to configure or remove prefix-list configuration on devices:
   You can also use this command to add or remove rules.

```
efa policy prefix-list update -type [ipv4|ipv6] --name [list name] --operation
[operation name]
```

Note

- The `add-device` and `remove-device` operations configure or remove a prefix list rules on the specified devices.
- The `add-rule` and `remove-rule` operations configure or remove a prefix list rules on the specified devices. If the prefix list is configured on the device, the rule is added or removed from the device.

The following is an example of IPv4 prefix list update:

- Add device

  The following example configures prefix list on the devices:

```
efa policy prefix-list update --name prefix_v4 --type ipv4 --operation add-device
--ip 10.20.246.10-11
```

- Delete device

  The following example removes prefix list from the devices:

```
efa policy prefix-list update --name prefix_v4 --type ipv4 --operation remove-
device --ip 10.20.246.10-11
```

- Add rule

  The following example adds rule to the already created prefix list:

```
efa policy prefix-list update --name prefix_v4 --type ipv4 --operation add-rule --
rule seq[5],action[permit],prefix[10.0.0.0/8],ge[16]
```

- Delete rule

  The following example removes rule from the existing prefix list:

```
efa policy prefix-list update --name prefix_v4 --type ipv4 --operation remove-rule
--rule seq[5],action[permit],prefix[10.0.0.0/8],ge[16]
```

The following example updates an IPv6 prefix list:

```
efa policy prefix-list update --type ipv6 --name prefix_1_in --operation add-device
--ip 10.20.246.29-30
+-------------+------+---------+--------+---------------+----+----+
|    Name     | Type | Seq num | Action |    Prefix     | Ge | Le |
+-------------+------+---------+--------+---------------+----+----+
| prefix_1_in | ipv6 | 11      | permit | 2001:db8::/48 |    |    |
+-------------+------+---------+--------+---------------+----+----+
Prefix-list details
+--------------+---------+--------+
|  IP Address  | Result  | Reason |
+--------------+---------+--------+
| 10.20.246.29 | Success |        |
+--------------+---------+--------+
| 10.20.246.30 | Success |        |
+--------------+---------+--------+
```

```
Device Results

efa policy prefix-list update --type ipv6 --name prefix_1_in --rule
"seq[11],action[permit],prefix[2001:db8::/48],ge[64],le[128]" --operation add-rule
+-------------+------+---------+--------+---------------+----+-----+
|    Name     | Type | Seq num | Action |    Prefix     | Ge | Le  |
+-------------+------+---------+--------+---------------+----+-----+
| prefix_1_in | ipv6 | 13      | permit | 2001:db8::/32 | 48 | 128 |
+-------------+------+---------+--------+---------------+----+-----+
| prefix_1_in | ipv6 | 14      | permit | 2003:db8::/32 | 64 | 128 |
+-------------+------+---------+--------+---------------+----+-----+
| prefix_1_in | ipv6 | 15      | deny   | 2003:db8::/63 | 64 | 128 |
+-------------+------+---------+--------+---------------+----+-----+
| prefix_1_in | ipv6 | 11      | permit | 2001:db8::/48 | 64 | 128 |
+-------------+------+---------+--------+---------------+----+-----+
Prefix-list details
+--------------+---------+--------+
|  IP Address  | Result  | Reason |
+--------------+---------+--------+
| 10.20.246.29 | Success |        |
+--------------+---------+--------+
| 10.20.246.30 | Success |        |
+--------------+---------+--------+
Device Results
```

a. Verify the switch configuration on the SLX device.

```
SLX# show running-config ip prefix-list
ip prefix-list prefix_v4 seq 5 permit 10.0.0.0/8 ge 16
```

3. Run the following command to show the IPv4 prefix list on a list of devices:

```
efa policy prefix-list list ?
Flags:
       --type string      Type of prefix-list. is ipv4 or ipv6
       --ip string        Comma separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2
```

The following example shows an IPv4 prefix list:

IPv4 prefix list show

```
efa policy prefix-list list --type ipv4 --ip 10.20.246.10-11

Name: prefix_v4
+------+---------+--------+------------+----+----+--------------+-------------+
| Type | Seq num | Action |   Prefix   | Ge | Le | DeviceIP     | AppState    |
+------+---------+--------+------------+----+----+--------------+-------------+
| ipv4 |    5    | permit | 10.0.0.0/8 | 16 |    | 10.20.246.10 | cfg-in-sync |
+------+---------+--------+------------+----+----+--------------+-------------+
| ipv4 |    5    | permit | 10.0.0.0/8 | 16 |    | 10.20.246.11 | cfg-in-sync |
+------+---------+--------+------------+----+----+--------------+-------------+
```

The following example shows an IPv6 prefix list:

IPv6 prefix list show

```
efa policy prefix-list list --type ipv6 --ip 10.20.246.29-30

Prefix-list details:

Name: prefix_1_in
+------+---------+--------+---------------+----+----+--------------+------------------+
| Type | Seq num | Action |    Prefix     | Ge | Le |   DeviceIP   |     AppState     |
+------+---------+--------+---------------+----+----+--------------+------------------+
```

```
| ipv6 | 11      | permit | 2001:db8::/48 |    |    | 10.20.246.29 | cfg-in-sync      |
+------+---------+--------+---------------+----+----+--------------+------------------+
| ipv6 | 11      | permit | 2001:db8::/48 |    |    | 10.20.246.30 | cfg-in-sync      |
+------+---------+--------+---------------+----+----+--------------+------------------+

efa policy prefix-list list --type ipv6
Prefix-list details:
Name: prefix_1_in
+------+---------+--------+---------------+----+----+
| Type | Seq num | Action |    Prefix     | Ge | Le |
+------+---------+--------+---------------+----+----+
| ipv6 | 11      | permit | 2001:db8::/48 |    |    |
+------+---------+--------+---------------+----+----+
```

4. Run the following command to delete the IPv4 prefix list on all devices:

   This step deletes the prefix list on all devices and XCO.

```
efa policy prefix-list delete ?
Flags:
  --type string      Type of prefix-list. is ipv4 or ipv6
  --name string      Name of Prefix list
```

   The following example deletes an IPv4 prefix list with name prefix_v4:

```
efa policy prefix-list delete --type ipv4 --name prefix_v4
```

   The following example deletes an IPv6 prefix list:

   System validates the IP prefix list name and type before running the delete operation. If a prefix list is bound to BGP peer or peer-group, an attempt to delete prefix-list will check for the presence of binding and report an error.

```
efa policy prefix-list delete --name plist2 --type ipv6
+--------+------+---------+
|  Name  | Type | Status  |
+--------+------+---------+
| plist2 | ipv6 | Success |
+--------+------+---------+
Prefix-list details
+--------------+---------+--------+
|  IP Address  | Result  | Reason |
+--------------+---------+--------+
| 10.20.246.29 | Success |        |
+--------------+---------+--------+
| 10.20.246.30 | Success |        |
+--------------+---------+--------+
```

Note
For more information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Drift and Reconcile (DRC) and Idempotency for IP Prefix List Configuration

The following table captures the various attributes of IP prefix list for which DRC and idempotency is supported.

- A drift is identified if any of the fields are modified through SLX, CLI or other management tools.

- A reconcile operation pushes the intended configuration to SLX, so bringing the SLX configuration in sync with XCO.

Field	Identify Drift	Reconcile config	Idempotency	Comments
IPv4 prefix list rule is deleted.	Yes	Yes	No	Deleted rule will be reconciled.
IPv4 prefix list is deleted.	Yes	Yes	No	Deleted prefix list along with all rules associated with it will be reconciled.
IPv4 prefix list rule created OOB. Different rules exist with same prefix list name in XCO.	No	No	NA	
IPv4 prefix-list rule created OOB. Different rules exist with same prefix-list name and same sequence number in XCO.	Yes	Yes	NA	Prefix list rule will be reconciled to be in sync with XCO.
Create a IPv4 prefix list OOB with a prefix list name not matching any of XCO created entries.	No	No	NA	These are treated as out of band entries and XCO will not perform DRC.

Drift and Reconcile (DRC) for IPv6 Prefix List

The drift and reconcile (DRC) for IPv6 prefix list is similar to IPv4 prefix list. If you associate a prefix list is with a route map whose association was changed on the device, then the DRC process computes the diff and reports the differences.

```
Config Drift: Prefix List
+-------------+------+-----+-------------------+--------------+
|    NAME     | TYPE | SEQ |     APP STATE     | CHILD CONFIG |
+-------------+------+-----+-------------------+--------------+
| prefix_2_in | ipv6 | 11  |  cfg-entry-added  |              |
| prefix_2_in | ipv6 | 12  |  cfg-entry-added  |              |
+-------------+------+-----+-------------------+--------------+

Config Drift: Route Map
+-----------+--------+-----+-------------------+--------------------------------+
|   NAME    | ACTION | SEQ |     APP STATE     |        CHILD CONFIG            |
+-----------+--------+-----+-------------------+--------------------------------+
| rmap_test | permit | 9   | cfg-entry-added   | matchIPv6PrefixList            |
|           |        |     |                   | prefix_3_in                    |
| rmap_test | permit | 9   | cfg-entry-deleted | matchIPv6PrefixList            |
|           |        |     |                   | prefix_2_in                    |
| rmap_test | permit | 10  | cfg-entry-deleted | matchIPv6PrefixList            |
|           |        |     |                   | prefix_3_in                    |
| rmap_test | permit | 10  | cfg-entry-added   | matchIPv6PrefixList            |
```

```
|           |        |     |                   | prefix_2_in                      |
+-----------+--------+-----+-------------------+----------------------------------+
```

Route Map

Route map is a route policy. It can use prefix list, access list, as-path, and community list to create an effective route policy. A route map consists of series of statements that check if a route matches the policy to permit or deny a route.

XCO enables you to create, delete, and update the route maps on a set of devices in fabric. Note that the IPv4 prefix list is the ONLY supported match criterion, and no other match criteria is supported. Also, no set criteria is supported.

Configure Route Map on devices

You can configure route map on a device.

About This Task

Follow this procedure to configure route map on devices.

Procedure

1. Run the following command to configure the route map with one or more rules:

```
efa policy route-map create ?
Flags:
      --name string           Name of route-map
      --rule stringArray      Rule in format seq[seq-num],action[permit/deny]
```

   The following is an example of creating a route map rmap_1 with two rules:

```
efa policy route-map create --name rmap_1 --rule seq[5],action[permit] --rule
seq[10],action[permit]
```

2. Run the following command to update the route map configuration on a list of devices:

   The update command configures the route map on device, removes configuration from a device or updates action of route-maps.

```
efa policy route-map update ?
Flags:
      --name string      Name of route-map
      --rule string      Rule in format seq[seq-num],action[permit/deny]
      --operation string Valid options are add-device, remove-device, update-action
      --ip string        Comma separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2
```

   - You can associate a route map with multiple rules.
   - The `add-device` operation adds all the rules of the route map on the specified devices.

- The `delete-device` operation deletes all the rules of the route map on the specified devices.

Note
The `delete-device` operation fails if the route map is bound to any BGP neighbor.

For an `update-action` operation, specify the route map name and the rule. You can modify the action to permit or deny for a specific rule. You can provide only one rule at a time.

The following are the examples of route map configuration update:

- **Add device:** Configures a route map rule on devices 10.20.246.10 and 10.20.246.11. Assume there are two route map rules for a map named rmap_1 that already exists in XCO:

  a. rmap_1 seq 5 action permit
  b. rmap_1 seq 10 action permit

```
efa policy route-map update --name rmap_1 --operation add-device --ip
10.20.246.10-11
```

- **Delete device:** Removes route map from the specified devices:

```
efa policy route-map update --name rmap_1 --operation delete-device --ip
10.20.246.10-11
```

- **Update action:** Changes the action from permit to deny for the specified rule:

```
efa policy route-map update --name rmap_1 --rule seq[5],action[deny] --operation
update-action
```

a. Verify the switch configuration on the SLX device.

Example 1	Example 2

`SLX# show running-config route-map`
`route-map rmap_1 permit 5`
	`SLX# show running-config route-map`
	`route-map rmap_1 permit 5`
	`route-map rmap_1 permit 10`

Example 3	

`SLX# show running-config route-map`
`route-map rmap_1 deny 5`
`route-map rmap_1 permit 10`	

3. Run the following command to create route map match criteria:

```
efa policy route-map-match create ?
Flags:
      --name string              Name of route-map
      --rule string              Rule in format seq[seq-num],action[permit/deny]
      --match-ipv6-prefix string IPv6 prefix-list name
```

The following is an example of route map match create in IPv6:

```
efa policy route-map-match create --name rmap_1 --rule seq[5],action[permit] --match-
ipv6-prefix prefix_1
```

a. Verify the switch configuration on the SLX device.

```
SLX# show running-config route-map
route-map rmap_1 permit 5
match ip address prefix-list prefix_1
```

4. Run the following command to remove the route map match criteria:

   The IPv6 prefix list is the only match supported.

```
efa policy route-map-match delete ?
Flags:
      --name string              Name of route-map
      --rule string              Rule in format seq[seq-num],action[permit/deny]
      --match-ipv6-prefix string IPv6 prefix-list name
```

   The following is an example of route map match delete in IPv6:

```
efa policy route-map-match delete --name rmap_1 --rule seq[5],action[permit]
```

   a. Verify the switch configuration on the SLX device.

```
SLX# show running-config route-map
route-map rmap_1 permit 5
```

5. Run the following command to display the route map for a list of devices:

   In the command output, the App State column reflects the state of configuration on the specified device. When there is drift in a rule, the App State is shown as `cfg-refreshed`.

```
efa policy route-map list ?
Flags:
      --ip string     Comma separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2
```

   Example:

```
efa policy route-map list --ip 10.20.246.10-11

Route-map details:

Name: rmap_1
Seq: 5
Action: permit
Match-ipv6-prefixlist:
    Prefix-list: prefix_1

Name: rmap_2
Seq: 5
Action: permit
Match-ipv6-prefixlist:
    Prefix-list: prefix_1

IP Addresses:
+----------------+-----+---------------+-----------------+
|     Name       | Seq |  IP Address   |     App State   |
+--------------+-----+---------------+-------------------+
| rmap_1       |  5  | 10.20.246.10  |    cfg-in-sync    |
+--------------+-----+---------------+-------------------+
| rmap_1       |  5  | 10.20.246.11  |    cfg-in-sync    |
+--------------+-----+---------------+-------------------+
| rmap_2       |  5  | 10.20.246.10  |    cfg-in-sync    |
+--------------+-----+---------------+-------------------+
| rmap_2       |  5  | 10.20.246.11  |    cfg-in-sync    |
+--------------+-----+---------------+-------------------+
```

6. Run the following command to delete a route map and the associated rules on the devices:

```
efa policy route-map delete ?
Flags:
```

```
      --name stringArray      Name of route-map
      --seq string            Sequence numbers. For example 5,10,20, or all
```

- The command removes the route map rule from the XCO database and from the associated devices.
- You can delete a specific rule of a route map by specifying the route map name and the sequence number of the rule.
- You can delete all the route map rules for a specific route map name by specifying the sequence number as "all".
- The result of this command depends on whether the route map is bound with a BGP neighbor.
  - If the route map is bound to BGP peer, you cannot delete the last route map rule.
  - If the route map has no bindings, the command deletes the configuration on all devices associated with the route map.

The following example deletes two rules with sequence numbers 5 and 10 from a route map (rmap_1) that has three rules:

- rmap_1 seq 5 action permit
- rmap_1 seq 10 action permit
- rmap_1 seq 20 action permit

```
efa policy route-map delete --name rmap_1 --seq 5,10
```

a. Verify the switch configuration on the SLX device.

```
SLX1# show running-config route-map rmap_1
route-map rmap_1 permit 20
```

Drift and Reconcile (DRC) and Idempotency for Route Map Configuration

The following table captures the various attributes of route map for which DRC and idempotency is supported.

- A drift is identified if any of the fields are modified through SLX, CLI, or other management tools.
- A reconcile operation pushes the intended configuration to SLX, so bringing the SLX configuration in sync with XCO.

Table 20: IP prefix list attributes supporting DRC and idempotency

Field	Identify Drift	Reconcile config	Idempotency	Comments
Route map deleted	Yes	Yes	No	Recreate the route map along with the match criteria during reconcile
Route map rule action updated	Yes	Yes	No	Reconcile the route map action (permit/deny) for that rule

Table 20: IP prefix list attributes supporting DRC and idempotency (continued)

Field	Identify Drift	Reconcile config	Idempotency	Comments
Update IPv4 prefix list name in match criteria	Yes	Yes	No	Reconcile the IPv4 prefix list name
IPv4 prefix list match criteria deleted	Yes	Yes	NA	Reconcile the match criteria for IPv4 prefix list
A different match criteria NOT supported by XCO is added through OOB	No	No	NA	
A set criteria NOT supported by XCO is added through OOB	No	No	NA	
Route map is created through OOB and this is not present/created by XCO	No	No	NA	

Event Handling for IP Prefix List

Event Handling for IP Prefix List and Route Map

Inventory service	Policy service
Inventory service maintains the IPv4 prefix list and route map information in its DB. As part of device update, diff is computed for these entities, and if there is a diff, an event is published.	• Policy service subscribes to the events from inventory service-related IP prefix list and route map and updates the app state of these entities. The entities for which the app state is in `cfg-refreshed` or `cfg-entry-deleted`, is reconciled as part of DRC. • To handle attribute level drift, DB maintains a bitmap to show exactly which attribute has drifted as part of DRC show output. • Policy service publishes events when you create, update or delete prefix lists and route maps.

Event Handling for IP Prefix List and Large Community List

Inventory service	Policy service
• Inventory service maintains the large community-list and route map information in its DB. When you update a device, diff is computed for these entities. If there is a diff, an event will be published. • Inventory service acknowledges the large community-list and route map events from policy service and update its DB accordingly.	• Policy service subscribes to the events from inventory service-related large community-list and route map, and updates the app-state of these entities. The OOB entries on the devices are stored in the DB, and marked as `cfg-not-managed`. The entities for which the app state is in `cfg-refreshed` or `cfg-deleted` (managed by XCO but changed on devices) state, are reconciled when DRC is done. • To handle attribute level drift, DB maintains a bitmap to show exactly which attribute has drifted according to DRC show command output. • Policy service publishes events when you create, update or delete community-list and route maps.

Community List

Use a community list to

- Create groups of BGP communities to use in a match part of a route map.
- Control which routes are accepted, preferred, distributed, or advertised.
- Set, append, or modify the communities of a route.

The following are the three types of community list which SLX supports.

1. **Standard:** 4 bytes BGP community
2. **Extended:** 8 bytes BGP community
3. **Large:** 12 bytes BGP community

XCO supports configuration of standard and extended community list. You can create, delete, update, and list standard and extended community list on a set of devices in fabric.

Create and update operations support rollback. If a configuration on one device fails, the configuration is rolled back on successfully configured devices. This ensures the consistent configuration of policies across all devices.

Delete operations do not support rollback. As a best practice, remove the configuration. If removing configuration on a device fails, the command displays an error for that device. The configuration that was removed successfully on other devices, will not be added back.

Configure Standard Community List

You can configure standard community list.

Before You Begin

- Ensure that the community list name begins with an alphabet followed by one or more alphanumeric characters.
- Ensure that the community list rule is inside single or double quotes.
- If the community list is not associated with a device, the created community rules are stored in XCO DB only. If the community list is already associated with a device, the created rules are also pushed to the devices in addition to stored in DB.

About This Task

Follow this procedure to configure standard community list.

Procedure

1. Run the following command to create a standard community list:

```
efa policy community-list create [flags]

Flags:
      --name string        Name of the community list.
      --type string        Type of the community list. Valid options are standard,
extended
      --rule stringArray   Rule in format seq[seq-num],action[permit/deny],std-
value[<1-4294967295>|<AA:NN, AA & NN is 2 bytes>|internet|local-as|no-export|no-
advertise] (or) ext-value[regular expression].

Example: "seq[5],action[permit],std-value[6550:125;local-as;internet]" (or)
"seq[4],action[deny],ext-value[^65000:.*_]"
```

Example:

```
efa policy community-list create --name comm1 --type standard --rule
"seq[5],action[permit],std-value[100;11:22;local-as;no-export]"

+-----------+-----+--------+--------------------+-------+
| Community | Seq | Action |     Std Value      | Ext   |
| List Name | num |        |                    | Value |
+-----------+-----+--------+--------------------+-------+
| comm-prye | 55  | permit | 100 11:22          |       |
|           |     |        | local-as no-export |       |
+-----------+-----+--------+--------------------+-------+
Community List details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results

efa policy community-list create --name stdext1 --type extended --rule
"seq[5],action[permit],ext-value[_2000_]"

+---------------------+---------+--------+-----------+-----------+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+---------------------+---------+--------+-----------+-----------+
| stdext1             | 5       | permit |           | _2000_    |
+---------------------+---------+--------+-----------+-----------+
Community List details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
```

```
+------------+--------+--------+-----------------+
Device Results
```

2. Run the following command to update a community list.

```
efa policy community-list update [flags]

Flags:
      --name string        Name of the community list.
      --type string        Type of the community list. Valid options are standard,
extended
      --rule string        Rule in format seq[seq-num],action[permit/deny],std-
value[<1-4294967295>|<AA:NN, AA & NN is 2 bytes>|internet|local-as|no-export|no-
advertise] (or) ext-value[regular expression]. Example: seq[5],action[permit],std-
value[6550:125;local-as;internet] (or) seq[4],action[deny],ext-value[^65000:.*_]
      --operation string   Valid options are update-rule, add-device, remove-device
      --ip string          Comma separated range of device IP addresses. Example:
"1.1.1.1-3","1.1.1.2","2.2.2.2"
```

Example:

- Add Device

```
efa policy community-list update --name stdext1 --type extended --operation add-
device --ip 10.20.246.29-30
+---------------------+---------+--------+-----------+-----------+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+---------------------+---------+--------+-----------+-----------+
| stdext1             | 4       | deny   |           | _1000_    |
+---------------------+---------+--------+-----------+-----------+
| stdext1             | 5       | permit |           | _2000_    |
+---------------------+---------+--------+-----------+-----------+
| stdext1             | 7       | deny   |           | _3000_    |
+---------------------+---------+--------+-----------+-----------+
Community List details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
+--------------+---------+--------+-----------------+
| 10.20.246.29 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.246.30 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results

show running-config ip community-list
ip community-list extended stdext1 seq 4 deny _1000_
ip community-list extended stdext1 seq 5 permit _2000_
ip community-list extended stdext1 seq 7 deny _3000_
```

Verify the switch configuration on SLX devices.

```
SLX# show running-config ip community-list
ip community-list standard comm1 seq 5 permit 0:100 11:22 local-as no-export
ip community-list extended commExt1 seq 3 permit _30000_
```

- Delete Device

```
efa policy community-list update --name comm1 --type standard --operation remove-
device --ip 10.20.63.140-141
+---------------------+---------+--------+-----------+-----------+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+---------------------+---------+--------+-----------+-----------+
| comm1               | 3       | permit | 65:12     |           |
+---------------------+---------+--------+-----------+-----------+
Community List details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
```

```
+--------------+---------+--------+-----------------+
| 10.20.63.140 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.63.141 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results
```

- Update rule

```
efa policy community-list update --name commExt1 --type extended --operation update-
rule --rule "seq[1],action[permit],ext-value[_30000_]"
+---------------------+---------+--------+
| Community List Name | Seq num | Action |
+---------------------+---------+--------+
| commExt1            | 1       | permit |
+---------------------+---------+--------+
Community List details
+---------------+----------+----------------------------------------+----------+
|  IP Address   |  Result  |                Reason                  | Rollback |
|               |          |                                        | reason   |
+---------------+----------+----------------------------------------+----------+
| 10.139.44.159 | Success  |                                        |          |
+---------------+----------+----------------------------------------+----------+
| 10.139.44.163 | Success  |                                        |          |
+---------------+----------+----------------------------------------+----------+
Device Results

On 10.139.44.159:
show running-config ip community-list
ip community-list extended commExt1 seq 30 action permit _30000_

efa policy community-list update --name comm1 --type standard --operation update-
rule "--rule seq[5]","action[permit]","std-value[100;no-advertise]"

+---------------+---------+-----------------------------------------+----------+
|  IP Address   |  Result |                  Reason                 | Rollback |
|               |         |                                         | reason   |
+---------------+---------+-----------------------------------------+----------|
| 10.139.44.159 | Failed  | Failed to create community list for     |          |
|               |         | comm1 on the device 10.139.44.159.      |          |
|               |         | Reason: For seq 5: netconf rpc [error]  |          |
|               |         | '%Error: Same filter is already configured |       |
|               |         | with sequencenumber 30.'                |          |
+---------------+---------+-----------------------------------------+----------+
| 10.139.44.163 | Rollback|                                         |          |
+---------------+---------+-----------------------------------------+----------+
Device Results
On 10.139.44.159:
show running-config ip community-list
ip community-list standard comm1 seq 30 action permit 100 no-advertise
```

3. Run the following command to delete a community list.

```
 efa policy community-list delete [flags]

Flags:
      --name string   Name of the community list.
      --type string   Type of the community list. Valid options are standard, extended.
      --seq string    Sequence numbers. For example 5,10,20 or all
```

- The CLI deletes the standard community list rules on all devices for the name, type, and sequence number provided and then deletes the community list rules from XCO.
- Pre-validation is done for seq IDs provided or for all sequence ids in case of 'all'. If any out-of-band, seq ID is provided in the request (or 'all' is specified and any

out-of-band seq ID exists), the operation is errored out without proceeding to remove config from device or XCO DB.

- You must either provide only XCO managed seq IDs in the CLI or REST request or remove the out-of-band seq IDs from device and execute the CLI or REST request again.

Example:

```
efa policy community-list delete --name commExt1 --seq all --type standard
+---------------------+---------+--------+
| Community List Name | Seq num | Action |
+---------------------+---------+--------+
| commExt1            | 1       | permit |
+---------------------+---------+--------+
| commExt1            | 2       | permit |
+---------------------+---------+--------+
| commExt1            | 3       | permit |
+---------------------+---------+--------+
Community List details
+---------------+---------+--------------+----------+
| IP Address    | Result  |    Reason    | Rollback |
|               |         |              | reason   |
+---------------+---------+--------------+----------+
| 10.139.44.159 | Success |              |          |
+---------------+---------+--------------+----------+
| 10.139.44.163 | Success |              |          |
+---------------+---------+--------------+----------+
Device Results
```

4. Run the following command to show a community list.

```
efa policy community-list list [flags]

Flags:
      --ip string     Comma separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2
      --name string   Name of the community list.
      --type string   Type of the community list. Valid options are standard, extended
```

Example,

```
efa policy community-list list

Community list details:

Name: clist1
Seq: 5
Action: deny
StdValue: 50:125 internet local-as no-advertise
ExtValue:

Name: clist1
Seq: 15
Action: deny
StdValue: 50:125 local-as
ExtValue:

Name: clist2
Seq: 1
Action: permit
StdValue:
ExtValue: _2000_

efa policy community-list list --type standard --ip 10.20.246.29-30
```

```
Community list details:

Name: clist1
Seq: 5
Action: deny
StdValue: 50:125 internet local-as no-advertise
ExtValue:

Name: clist1
Seq: 15
Action: deny
StdValue: 50:125 local-as
ExtValue:

IP Addresses:
+--------+-----+--------------+-------------+
|  Name  | Seq |  IP Address  |  App State  |
+--------+-----+--------------+-------------+
| clist1 | 5   | 10.20.246.29 | cfg-in-sync |
+--------+-----+--------------+-------------+
| clist1 | 5   | 10.20.246.30 | cfg-in-sync |
+--------+-----+--------------+-------------+
| clist1 | 15  | 10.20.246.29 | cfg-in-sync |
+--------+-----+--------------+-------------+
| clist1 | 15  | 10.20.246.30 | cfg-in-sync |
+--------+-----+--------------+-------------+
```

Rollback Support

A new column “Rollback reason” in the device results output displays the reason when the rollback operation has failed.

If the “Result” column displays “Failed”, the “Reason” and “Rollback reason” columns display sufficient information to capture why the operation has failed.

If the “Result” column displays “Rollback”, then the given operation has been rollbacked successfully on the associated device.

Default value of “Success” means everything was OK on that device for the given operation.

```
efa policy community-list create --name commExt1 --type extended --rule
"seq[3],action[permit],ext-value[_30000_]"

+---------------------+---------+--------+
| Community List Name | Seq num | Action |
+---------------------+---------+--------+
| commExt1            | 3       | permit |
+---------------------+---------+--------+
Community List details
+---------------+----------+-----------------------------------------+-----------+
|  IP Address   |  Result  |                 Reason                  | Rollback  |
|               |          |                                         | reason    |
+---------------+----------+-----------------------------------------+-----------+
| 10.139.44.159 | Failed   | Failed to create community list for     |           |
|               |          | commExt1 on the device 10.139.44.159.   |           |
|               |          | Reason: For seq 3: netconf rpc [error]  |           |
|               |          | '%Error: Same filter is already         |           |
|               |          | configured with sequence number 30.'    |           |
+---------------+----------+-----------------------------------------+-----------+
```

```
| 10.139.44.163 | Rollback |                                                |          |
+---------------+----------+------------------------------------------------+----------+
Device Results
```

Configure Extended Community List

You can configure an extended community list.

Before You Begin

- Ensure that the extended community list (extcommunity-list) name begins with an alphabet followed by one or more alphanumeric characters.
- Ensure that the extended community list rule is inside single or double quotes.
- If the extended community list is not associated with a device, the created community rules are stored in XCO DB only. If the excommunity list is already associated with a device, the created community rules are also pushed to the devices and stored in XCO DB.

For supported commands on extended community list, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

About This Task

Follow this procedure to configure an extended community list.

Procedure

1. Run the following command to create an extended community list:

```
efa policy extcommunity-list create
```

The following is an example of configuring an extended community list:

```
efa policy extcommunity-list create --name excommlist-1 --type standard --rule
"seq[4],action[permit],soo[10.11.2.3:22]"
efa policy extcommunity-list create --name excommlist-1 --type standard --rule
"seq[5],action[deny],rt[1:345]"
efa policy extcommunity-list create --name excommlist-1 --type standard --rule
"seq[6],action[permit],rt[1:45],soo[10.11.2.3:22]"
efa policy extcommunity-list create --name excommlist-1 --type standard --rule
"seq[7],action[deny],rt[1:345],soo[6:12]"

efa policy extcommunity-list create --name excommlist-2 --type extended --rule
"seq[2],action[permit],ext-value[_15000_]"
efa policy extcommunity-list create --name excommlist-2 --type extended --rule
"seq[5],action[deny],ext-value[_20000_]"

+--------------------+-----+--------+----+--------------+-------+
| Extended community | Seq | Action | Rt |     Soo      | Ext   |
| list name          | num |        |    |              | Value |
+--------------------+-----+--------+----+--------------+-------+
| excommlist-1       | 4   | permit |    | 10.11.2.3:22 |       |
+--------------------+-----+--------+----+--------------+-------+
Extended community list details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results

+--------------------+-----+--------+-------+--------------+-------+
| Extended community | Seq | Action |  Rt   |     Soo      | Ext   |
```

```
| list name          | num |        |       |              | Value |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 5   | deny   | 1:345 |              |       |
+--------------------+-----+--------+-------+--------------+-------+
Extended community list details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results

+--------------------+-----+--------+------+--------------+-------+
| Extended community | Seq | Action |  Rt  |     Soo      | Ext   |
| list name          | num |        |      |              | Value |
+--------------------+-----+--------+------+--------------+-------+
| excommlist-1       | 6   | permit | 1:45 | 10.11.2.3:22 |       |
+--------------------+-----+--------+------+--------------+-------+
Extended community list details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results

+--------------------+-----+--------+----+-------+---------+
| Extended community | Seq | Action | Rt |  Soo  | Ext     |
| list name          | num |        |    |       | Value   |
+--------------------+-----+--------+----+-------+---------+
| excommlist-2       | 2   | permit |    |       | _15000_ |
+--------------------+-----+--------+----+-------+---------+
Extended community list details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results

+--------------------+-----+--------+----+-----+---------+
| Extended community | Seq | Action | Rt | Soo | Ext     |
| list name          | num |        |    |     | Value   |
+--------------------+-----+--------+----+-----+---------+
| excommlist-2       | 5   | deny   |    |     | _20000_ |
+--------------------+-----+--------+----+-----+---------+
Extended community list details

+--------------------+-----+--------+---+-------------+------+
| Extended community | Seq | Action | Rt|     Soo     | Ext  |
| list name          | num |        |   |             | Value|
+--------------------+-----+--------+---+-------------+------+
| excommlist-1       | 4   | permit |   |10.11.2.3:22 |      |
+--------------------+-----+--------+---+-------------+------+
Extended community list details

+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results
```

2. Run the following command to update an extended community list:

```
efa policy extcommunity-list update
```

Example:

- The following is an example of adding a device when you update an extended community list:

```
efa policy extcommunity-list update --name excommlist-1 --type standard --operation
add-device --ip 10.20.246.29,10.20.246.30
```

```
+--------------------+-----+--------+-------+--------------+-------+
| Extended community | Seq | Action |  Rt   |     Soo      | Ext   |
| list name          | num |        |       |              | Value |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 4   | permit |       | 10.11.2.3:22 |       |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 5   | deny   | 1:345 |              |       |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 6   | permit | 1:45  | 10.11.2.3:22 |       |
+--------------------+-----+--------+-------+--------------+-------+
Extended community list details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
+--------------+---------+--------+-----------------+
| 10.20.246.29 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.246.30 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results


efa policy extcommunity-list update --name excommlist-2 --type extended --operation
add-device --ip 10.20.246.29,10.20.246.30
+------------------------------+---------+--------+----+-----+-----------+
| Extended community list name | Seq num | Action | Rt | Soo | Ext Value |
+------------------------------+---------+--------+----+-----+-----------+
| excommlist-2                 | 2       | permit |    |     | _15000_   |
+------------------------------+---------+--------+----+-----+-----------+
| excommlist-2                 | 5       | deny   |    |     | _20000_   |
+------------------------------+---------+--------+----+-----+-----------+
Extended community list details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
+--------------+---------+--------+-----------------+
| 10.20.246.29 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.246.30 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results
```

- The following is an example of deleting a device when you update an extended community list:

```
efa policy extcommunity-list update --name excommlist-2 --type extended --operation
remove-device --ip 10.20.246.29,10.20.246.30
+--------------------+-----+--------+----+-----+----------+
| Extended community | Seq | Action | Rt | Soo | Ext      |
| list name          | num |        |    |     | Value    |
+--------------------+-----+--------+----+-----+----------+
| excommlist-2       | 2   | permit |    |     | _15000_  |
+--------------------+-----+--------+----+-----+----------+
| excommlist-2       | 5   | deny   |    |     | _25000_  |
+--------------------+-----+--------+----+-----+----------+
Extended community list details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
+--------------+---------+--------+-----------------+
| 10.20.246.29 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.246.30 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results

efa policy extcommunity-list update --name excommlist-1 --type standard --operation
```

```
remove-device --ip 10.20.246.29,10.20.246.30

+--------------------+-----+--------+-------+--------------+-------+
| Extended community | Seq | Action |  Rt   |     Soo      | Ext   |
| list name          | num |        |       |              | Value |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 4   | permit |       | 10.11.2.3:22 |       |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 5   | deny   | 1:345 |              |       |
+--------------------+-----+--------+-------+--------------+-------+
| excommlist-1       | 6   | permit | 1:45  | 10.11.2.3:22 |       |
+--------------------+-----+--------+-------+--------------+-------+
Extended community list details
+--------------+--------+--------------------------------------+----------+
|  IP Address  | Result |                Reason                | Rollback |
|              |        |                                     | reason   |
+--------------+--------+--------------------------------------+----------+
| 10.20.246.29 | Failed | Device 10.20.246.29 not reachable. |          |
|              |        | Please retry after verifying the   |          |
|              |        | inputs and connectivity issues     |          |
+--------------+--------+--------------------------------------+----------+
| 10.20.246.30 | Failed | Device 10.20.246.30 not reachable. |          |
|              |        | Please retry after verifying the   |          |
|              |        | inputs and connectivity issues     |          |
+--------------+--------+--------------------------------------+----------+
Device Results
```

- The following is an example of updating a rule when you update an extended community list:

```
efa policy extcommunity-list update --name excommlist-2 --type extended --operation
update-rule --rule "seq[5],action[deny],ext-value[_25000_]"

+------------------------------+---------+--------+----+-----+-----------+
| Extended community list name | Seq num | Action | Rt | Soo | Ext Value |
+------------------------------+---------+--------+----+-----+-----------+
| excommlist-2                 | 5       | deny   |    |     | _25000_   |
+------------------------------+---------+--------+----+-----+-----------+
Extended community list details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
+--------------+---------+--------+-----------------+
| 10.20.246.29 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.246.30 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results

efa policy extcommunity-list update --name excommlist-1 --type standard --operation
update-rule --rule "seq[5],action[permit],rt[0:123],soo[0:12]"

+------------------------------+---------+--------+-------+------+-----------+
| Extended community list name | Seq num | Action |  Rt   | Soo  | Ext Value |
+------------------------------+---------+--------+-------+------+-----------+
| excommlist-1                 | 5       | permit | 0:123 | 0:12 |           |
+------------------------------+---------+--------+-------+------+-----------+
Extended community list details
+--------------+--------+-----------------------------------------------+---------+
|  IP Address  | Result |                  Reason                       | Rollback|
|              |        |                                               | reason  |
+--------------+--------+-----------------------------------------------+---------+
| 10.20.246.29 | Failed |Reason: For seq 5: netconf rpc [error]    |         |
|              |        |'"rt 0:123 soo 0:12" is an invalid value.'|         |
+--------------+--------+-----------------------------------------------+---------+
| 10.20.246.30 | Failed |Reason: For seq 5: netconf rpc [error]    |         |
```

```
|              |        |'"rt 0:123 soo 0:12" is an invalid value.'|         |
+--------------+--------+------------------------------------------+---------+
Device Results
```

3. Run the following command to delete an extended community list:

```
efa policy extcommunity-list delete
```

   - The CLI deletes the extended community list rules on all devices for the name, type, and sequence provided and then deletes the extended community list rules from XCO.
   - Pre-validation is done for seq IDs provided or for all sequence IDs in case of 'all'. If any out-of-band and seq ID is provided in the request (or 'all' is specified and any out-of-band seq ID exists), the operation is errored out without proceeding to remove config from device or XCO DB.
   - You must either provide only XCO managed seq IDs in the CLI or REST request or remove the out-of-band seq IDs from device, and then run the CLI or REST request again.

   The following example deletes an extended community list:

```
efa policy extcommunity-list delete --name excommlist-2 --type extended --seq all

+------------------------------+--------+--------+----+----+----------+
| Extended community list name | Seq num| Action | Rt |Soo | Ext Value|
+------------------------------+--------+--------+----+----+----------+
| excommlist-2                 | 2      | permit |    |    | _15000_  |
+------------------------------+--------+--------+----+----+----------+
| excommlist-2                 | 5      | deny   |    |    | _25000_  |
+------------------------------+--------+--------+----+----+----------+
Extended community list details
+--------------+---------+--------+-----------------+
|  IP Address  | Result  | Reason | Rollback reason |
+--------------+---------+--------+-----------------+
| 10.20.246.29 | Success |        |                 |
+--------------+---------+--------+-----------------+
| 10.20.246.30 | Success |        |                 |
+--------------+---------+--------+-----------------+
Device Results
```

4. Run the following command to list the extended community-list on a list of devices or to filter by name or by type:

```
efa policy extcommunity-list list
```

   The following example shows an extended community list configuration on list of devices:

```
efa policy extcommunity-list list

Extended community list details:


Name: excommlist-1
Seq: 5
Action: permit
Route Target: 1:100 2:200 3:145 4:123
Site of Origin: 1.2.3.4:12 5:400 10.11.12.13:22
ExtValue:

Name: excommlist-1
Seq: 6
Action: permit
Route Target: 1:45
```

```
Site of Origin: 10.11.2.3:22
ExtValue:

Name: excommlist-1
Seq: 9
Action: deny
Route Target: 1:345
Site of Origin: 6:12
ExtValue:

Name: excommlist-2
Seq: 2
Action: permit
Route Target:
Site of Origin:
ExtValue: _15000_


efa policy extcommunity-list list --ip 10.20.246.29 --name excommlist-1

Extended community list details:

Name: excommlist-1
Seq: 6
Action: permit
Route Target: 1:45
Site of Origin: 10.11.2.3:22
ExtValue:

Name: excommlist-1
Seq: 9
Action: deny
Route Target: 1:345
Site of Origin: 6:12
ExtValue:

IP Addresses:
+--------------+-----+--------------+-------------+
|     Name     | Seq |  IP Address  |  App State  |
+--------------+-----+--------------+-------------+
| excommlist-1 | 6   | 10.20.246.29 | cfg-in-sync |
+--------------+-----+--------------+-------------+
| excommlist-1 | 9   | 10.20.246.29 | cfg-in-sync |
+--------------+-----+--------------+-------------+

efa policy extcommunity-list list --ip 10.20.246.29 --name excommlist-2

Extended community list details:

Name: excommlist-2
Seq: 2
Action: permit
Route Target:
Site of Origin:
ExtValue: _15000_

Name: excommlist-2
Seq: 5
Action: deny
Route Target:
Site of Origin:
ExtValue: _20000_

IP Addresses:
```

```
+--------------+-----+--------------+-------------+
|     Name     | Seq |  IP Address  |  App State  |
+--------------+-----+--------------+-------------+
| excommlist-2 | 2   | 10.20.246.29 | cfg-in-sync |
+--------------+-----+--------------+-------------+
| excommlist-2 | 5   | 10.20.246.29 | cfg-in-sync |
+--------------+-----+--------------+-------------+
```

Configure Large Community List

You can configure a large community list.

About This Task

Follow this procedure to configure a large community list.

- If the large community list not associated with a device, the configured large community rules will be stored in the Policy service DB.
- If the large community list is already associated with devices, the configured large community rules will also be pushed to the devices.
- The large community list configuration supports rollback. The rollback will be attempted on all the associated devices.

For supported commands on large community list, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Procedure

1. Run the following command to create a large community list:

```
efa policy large-community-list create
```

Example:

```
efa policy large-community-list create --name lgcomm1 --type standard --rule
"seq[5],action[permit],std-value[10:10:10;20:20:20]"

+---------------------+---------+--------+-------------------+-----------+
| Community List Name | Seq num | Action |     Std Value     | Ext Value |
+---------------------+---------+--------+-------------------+-----------+
| lgcomm1             | 5       | permit | 10:10:10 20:20:20 |           |
+---------------------+---------+--------+-------------------+-----------+
Community List details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results

efa policy large-community-list create --name lgcommExt1 --type extended --rule
"seq[5],action[permit],ext-value[_645XX:.*:.*]"

+---------------------+---------+--------+-------------------+---------------- +
| Community List Name | Seq num | Action |     Std Value     |    Ext Value   |
+---------------------+---------+--------+-------------------+-----------------+
| lgcommExt1          | 5       | permit |                   |  _645XX:.*:.* |
+---------------------+---------+--------+-------------------+---------------- +
Community List details
+------------+--------+--------+-----------------+
| IP Address | Result | Reason | Rollback reason |
+------------+--------+--------+-----------------+
Device Results
```

```
efa policy large-community-list create --name lgcomm1 --type standard --rule
"seq[15],action[permit],std-value[10:10:10 20:20:20]"
+---------------------------+---------+--------+-------------------+-----------+
| Large Community List Name | Seq num | Action |     Std Value     | Ext Value |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 15      | permit | 10:10:10 20:20:20 |           |
+---------------------------+---------+--------+-------------------+-----------+
Large Community List details
+---------------+----------
+-------------------------------------------------------------------+-----------------+
|  IP Address   |  Result
|                            Reason                                 | Rollback reason |
+---------------+----------
+-------------------------------------------------------------------+-----------------+
| 10.139.44.159 | Failed
| Policy lgcomm1 type large-community-list seq# 15 operation   |                 |
|               |          |
failed on device 10.139.44.159 due to Reason: For seq        |                 |
|               |          |
15: netconf rpc [error] '%Error: Same filter is already      |                 |
|               |          |
configured with sequence number 10.'                         |                 |
+---------------+----------
+-------------------------------------------------------------------+-----------------+
| 10.139.44.160 | Rollback
|                                                                   |                 |
+---------------+----------
+-------------------------------------------------------------------+-----------------+
Device Results
```

2. Run the following command to update a large community list:

```
efa policy large-community-list update
```

You can use the **efa policy large-community-list update** command to update (add or remove) devices. Use the update operation to configure or deconfigure the large community list rules on a device or list of devices. The update operation supports rollback for add device where rollback is attempted on failed devices.

Example:

- The following is an example of adding a device when you update a large community list:

```
efa policy large-community-list update --operation add-device --ip 10.139.44.159 --
name lgcomm1 --type standard
+---------------------------+---------+--------+-------------------+-----------+
| Large Community List Name | Seq num | Action |     Std Value     | Ext Value |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 5       | permit | 10:10:10 20:20:20 |           |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 10      | permit | 30:30:30          |           |
+---------------------------+---------+--------+-------------------+-----------+
Community List details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results


efa policy large-community-list update --name lgcommExt1 --type extended --
operation add-device --ip 10.139.44.159
```

```
+---------------------+---------+--------+----------------------+-----------------+
| Community List Name | Seq num | Action |      Std Value       |    Ext Value.   |
+---------------------+---------+--------+----------------------+-----------------+
| lgcommExt1          | 5       | permit |                      |   645XX:.*:.*  |
+---------------------+---------+--------+----------------------+-----------------+
Community List details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

The following is an example of a switch configuration on SLX devices:

```
SLX# show running-config ip large-community-list
ip large-community-list standard lgcomm1 seq 5 permit 10:10:10 20:20:20
ip large-community-list standard lgcomm1 seq 10 permit 30:30:30
ip large-community-list extended lgcommExt1 seq 5 permit _645XX:.*:.*


efa policy large-community-list update --operation add-device --ip
10.139.44.159-160 --name lgcomm1 --type standard
+---------------------------+---------+--------+-------------------+-----------+
| Large Community List Name | Seq num | Action |     Std Value     | Ext Value |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 5       | permit | 10:10:10 20:20:20 |           |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 10      | permit | 30:30:30          |           |
+---------------------------+---------+--------+-------------------+-----------+
Community List details
+---------------+---------
+-------------------------------------------------------------------+-----------------+
|  IP Address   | Result
|                            Reason                                 | Rollback reason |
+---------------+---------
+-------------------------------------------------------------------+-----------------+
| 10.139.44.159 | Success
|                                                                   |                 |
+---------------+---------
+-------------------------------------------------------------------+-----------------+
| 10.139.44.160 | Failed
| Policy lgcomm1 type large-community-list seq# 5 operation    |                 |
|               |         |
failed on device 10.139.44.160 due to Reason: For seq 10:    |                 |
|               |         |
netconf rpc [error] '%Error: An IP Community access-list     |                 |
|               |         |
with this name and instance number already exists'           |                 |
+---------------+---------
+-------------------------------------------------------------------+-----------------+
Device Results
```

- The following is an example of removing a device when you update a large community list:

```
efa policy large-community-list update --name lgcomm1 --type standard --operation
remove-device --ip 10.139.44.159

+---------------------------+---------+--------+-------------------+-----------+
| Large Community List Name | Seq num | Action |     Std Value     | Ext Value |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 5       | permit | 10:10:10 20:20:20 |           |
+---------------------------+---------+--------+-------------------+-----------+
| lgcomm1                   | 10      | permit | 30:30:30          |           |
```

```
+--------------------------+---------+--------+------------------+-----------+
Community List details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

3. Run the following command to delete a large community list:

   The CLI deletes a large community list rules on all devices for the given type and sequence and then deletes the large community list rules from XCO.

```
efa policy large-community-list delete
```

   The following example deletes a large community list:

```
efa policy large-community-list delete --name lgcomm1 --seq all --type standard

+--------------------+---------+--------+------------------+-----------+
| Community List Name | Seq num | Action |     Std Value     | Ext Value |
+--------------------+---------+--------+------------------+-----------+
| lgcomm1             | 5       | deny   | 10:10:10 20:20:20 |           |
+--------------------+---------+--------+------------------+-----------+
| lgcomm1             | 10      | permit | 30:30:30          |           |
+--------------------+---------+--------+------------------+-----------+
Community List details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

4. Run the following command to list the large community list for a list of devices or to filter by name or by type:

```
efa policy large-community-list list
```

   The following example shows the large community list configuration on list of devices:

```
efa policy large-community-list list

large community list details:

Name: lgcomm1
Seq: 5
Action: deny
StdValue: 10:10:10 20:20:20
ExtValue:

Name: lgcomm1
Seq: 10
Action: permit
StdValue: 30:30:30
ExtValue:

Name: lgcommExt1
Seq: 5
Action: permit
StdValue:
ExtValue: _645XX:.*:.*

efa policy large-community-list list --type standard --ip 10.139.44.159
```

```
Large community list details:

Name: lgcomm1
Seq: 5
Action: deny
StdValue: 10:10:10 20:20:20
ExtValue:

Name: lgcomm1
Seq: 10
Action: permit
StdValue: 30:30:30
ExtValue:

IP Addresses:
+------------+-----+---------------+-----------------+
|    Name    | Seq |  IP Address   |    App State    |
+------------+-----+---------------+-----------------+
| lgcomm1    | 5   | 10.139.44.159 | cfg-in-sync     |
+------------+-----+---------------+-----------------+
| lgcomm1    | 10  | 10.139.44.159 | cfg-in-sync     |
+------------+-----+---------------+-----------------+
```

Drift and Reconcile (DRC), Idempotency for Standard and Extended Community-list Configuration

The following table describes the various attributes of standard and extended community list for which DRC and idempotency is supported. A drift is identified if any of the fields below are modified by user through SLX CLI or other management tools. Reconcile operation pushes the intended configuration to SLX, so bringing the SLX configuration in sync with XCO.

If you create an entry for idempotency which already exists in XCO, the system shows an error message stating that entry already exists.

Field	Identify Drift	Reconcile config	Idempotency	Comments
community-list rule is deleted	Yes	Yes	No	Deleted rule will be reconciled.
community-list is deleted	Yes	Yes	No	Deleted community-list along with all rules associated with it will be reconciled.
community-list rule is changed	Yes	Yes	No	Updated rule will be reconciled.
community-list rule created by OOB. Different rules exist with same community-list name in XCO	No	No	NA	

Field	Identify Drift	Reconcile config	Idempotency	Comments
Create a community-list OOB with a community-list name not matching any of XCO created entries	No	No	NA	These are treated as out of band entries and XCO will not perform DRC
extcommunity-list rule is deleted	Yes	Yes	No	Deleted rule will be reconciled.
extcommunity-list is deleted	Yes	Yes	No	Deleted community-list along with all rules associated with it will be reconciled.
extcommunity-list rule is changed	Yes	Yes	No	Updated rule will be reconciled.
extcommunity-list rule created by OOB. Different rules exist with same extcommunity-list name in XCO	No	No	NA	
Create an extcommunity-list OOB with a extcommunity-list name not matching any of XCO created entries	No	No	NA	These are treated as out of band entries and XCO will not perform DRC

Route Map Match and Set of Community List

Route map is a route policy which can use prefix list, access-lists, as-path, and community list to create a route policy. A route map consists of series of statements that check if a route matches the policy to permit or deny a route. You can also use set criteria to alter the properties of route as they are installed in the routing table.

In EFA 3.1.0, you can create or delete route map match and set criteria on a set of devices in fabric.

Create and update operations support rollback. If configuration on one device fails, the configuration is rolled back on devices that were successfully configured. This ensures consistent configuration of policies across all devices.

Delete operations do not support rollback. As a best practice, you must remove the configuration. If removing configuration on a device fails, the command displays an error for that device. The configuration which was removed successfully on other devices, will not be added back.

XCO supports CLIs for the following route map match and set criteria:

1. Route map match criteria to standard and extended community list
2. Route map set criteria of standard and extended community attributes

3. Route map set criteria of standard community-list delete

Configure Route Map Match and Set of Community List

You can configure route map match and set of community list.

About This Task

Follow this procedure to configure route map match and set of community list.

Procedure

1. Run the following command to create a route map match of a certain match type.

```
efa policy route-map-match create [flags]

Flags:
      --name string                      Name of the route map
      --rule string                      Rule in format seq[seq-num],action[permit/
deny]. Example: seq[5],action[permit]
      --match-ipv4-prefix string         IPv4 prefix-list name
      --match-community-list string      Community list name
      --match-extcommunity-list string   ExtCommunity list name
```

Example:

```
efa policy route-map-match create --name foo --rule seq[10],action[permit] --match-
extcommunity-list extFoo
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| foo            | 10      | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Success |        |                 |
+---------------+---------+--------+-----------------+
```

a. Verify the switch configuration on SLX devices.

```
route-map foo permit 10
match extcommunity extFoo
```

2. Run the following command to remove a route map match of a certain match type.

```
 efa policy route-map-match delete [flags]

Flags:
      --name string                      Name of the route map
      --rule string                      Rule in format seq[seq-num],action[permit/
deny]. Example: seq[5],action[permit]
      --match-ipv4-prefix string         IPv4 prefix-list name
      --match-community-list string      Community list name
      --match-extcommunity-list string   ExtCommunity list name
```

Example:

```
efa policy route-map-match delete --name foo --rule seq[10],action[permit] --match-
extcommunity-list extFoo
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| foo            | 10      | permit |
+----------------+---------+--------+
Route Map details
```

```
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Success |        |                 |
+---------------+---------+--------+-----------------+
```

a. Verify the switch configuration on SLX devices.

```
route-map foo permit 10
!
```

3. Run the following command to set the standard and extended community list attributes.

   The CLI sets the community list for deletion.

```
efa policy route-map-set create [flags]

Flags:
      --name string                       Name of the route map
      --rule string                       Rule in format seq[seq-num],action[permit/
deny]. Example: seq[5],action[permit]
      --set-community string              --set-community [<1-4294967295>|<AA:NN, AA
& NN is 2 bytes>|internet|local-as|no-export|no-advertise]. Example: 6550:125,local-
as,internet
      --set-extcommunity-rt string        --set-extcommunity-rt [ASN:NN|IpAddress:NN,
ASN & NN is 2 or 4 bytes | additive]. Example: 2:300,12.12.13.33:24
      --set-extcommunity-soo string       --set-extcommunity-soo [ASN:NN|IpAddress:NN,
ASN & NN is 2 or 4 bytes]. Example: 32:124
      --set-communitylist-delete string   --set-communitylist-delete [community-list
name]
```

   Example:

```
efa policy route-map-set create --name foo --rule seq[10],action[permit] --set-
community 6550:125,internet,local-as
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| foo            | 10      | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.161 | Success |        |                 |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
--- Time Elapsed: 10.33886575s ---
```

a. Verify the switch configuration on SLX devices.

```
SLX# show running-config route-map
route-map foo permit 10
set community 6550:125 local-as internet
!
```

4. Run the following command to delete the set directive on route map.

```
efa policy route-map-set delete [flags]

Flags:
      --name string                       Name of the route map
      --rule string                       Rule in format seq[seq-num],action[permit/
deny]. Example: seq[5],action[permit]
```

```
      --set-community string              --set-community [<1-4294967295>|<AA:NN, AA
& NN is 2 bytes>|internet|local-as|no-export|no-advertise]. Example: 6550:125,local-
as,internet
      --set-extcommunity-rt string        --set-extcommunity-rt [ASN:NN|IpAddress:NN,
ASN & NN is 2 or 4 bytes | additive]. Example: 2:300,12.12.13.33:24
      --set-extcommunity-soo string       --set-extcommunity-soo [ASN:NN|IpAddress:NN,
ASN & NN is 2 or 4 bytes]. Example: 32:124
      --set-communitylist-delete string   --set-communitylist-delete [community-list
name]
```

Example:

```
efa policy route-map-set delete --name foo --rule seq[10],action[permit] --set-
community 6550:125,internet,local-as
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| foo            | 10      | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.161 | Success |        |                 |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
--- Time Elapsed: 11.547377938s ---
```

a. Verify the switch configuration on SLX devices.

```
SLX# show running-config route-map
route-map foo permit 10
!
```

Drift and Reconcile (DRC) and Idempotency for Route Map Match and Set Configuration

The following table describes the various attributes of route map match and set for which DRC and idempotency is supported. A drift is identified if any of the fields mentioned in the following table are modified by user through SLX CLI or other management tools. Reconcile operation pushes the intended configuration to SLX, so bringing the SLX configuration in sync with XCO.

Field	Identify Drift	Reconcile config	Idempotency	Comments
Update community-list name in match criteria	Yes	Yes	No	Reconcile the community-list name
Community-list match criteria deleted	Yes	Yes	NA	Reconcile the match criteria for community-list
Update extcommunity-list name in match criteria	Yes	Yes	No	Reconcile the extcommunity-list name
Extcommunity-list match criteria deleted	Yes	Yes	NA	Reconcile the match criteria for extcommunity-list

Field	Identify Drift	Reconcile config	Idempotency	Comments
Update community-list name in set criteria	Yes	Yes	No	Reconcile the community-list name
Community-list set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for community-list
Update community attribute in set criteria	Yes	Yes	No	Reconcile the community attribute
Community attribute set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for community attribute
Update extcommunity rt attribute in set criteria	Yes	Yes	No	Reconcile the extcommunity rt attribute
Extcommunity rt attribute set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for extcommunity rt attribute
Update extcommunity soo attribute in set criteria	Yes	Yes	No	Reconcile the extcommunity soo attribute
Extcommunity soo attribute set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for extcommunity soo attribute
A different match criteria NOT supported by XCO is added through OOB	No	No	NA	These are treated as out of band entries and XCO will not store them
A different set criteria NOT supported by XCO is added through OOB	No	No	NA	These are treated as out of band entries and XCO will not store them
Route-map match criteria is created through OOB and this is not present/created by XCO	No	No	NA	These are treated as out of band entries and XCO will not perform DRC
Route-map set criteria is created through OOB and this is not present/created by XCO	No	No	NA	These are treated as out of band entries and XCO will not perform DRC

Route Map Match and Set of Large Community List

Route map is a route policy. It can use prefix-list, access-lists, as-path, large community list etc. to create an effective route policy. A route-map consists of series of statements that check if a route matches the policy to permit or deny a route. Also, set criteria can be used to alter the properties of route as they are installed in the routing table.

XCO 3.2.0 supports match and set criteria for large community list and IPv6 prefix list.

XCO supports the following CLIs:

- route-map match on large community list
- route-map set large community attributes
- route-map set large community-list delete

For supported commands on route map match and set, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Configure Route Map Match

You can configure a route map match.

About This Task

Follow this procedure to configure a route map match.

Procedure

1. Run the following command to create a route map match:

```
efa policy route-map-match create
```

Example:

```
efa policy route-map-match create --name rmap1 --rule "seq[1],action[permit]" --match-
largecommunity-list lgcomm1
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| rmap1          | 1       | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

2. Verify the following configuration on SLX devices:

```
SLX# show running-config route-map
route-map rmap1 permit 1
  match large-community-list lgcomm1
```

3. Run the following command to delete a route map match:

```
efa policy route-map-match delete
```

Example:

```
efa policy route-map-match delete --name rmap1 --rule "seq[1],action[permit]" --match-
largecommunity-list lgcomm1
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| rmap1          | 1       | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
```

```
+---------------+---------+--------+-----------------+
Device Results
```

4. Verify the following configuration on SLX devices:

```
SLX# show running-config route-map
route-map rmap1 permit 1
```

Configure Route Map Set

You can configure a route map set.

About This Task

Follow this procedure to configure a route map set.

Procedure

1. Run the following command to create a route map set:

```
efa policy route-map-set create
```

Example:

```
efa policy route-map-set create --name rmap1 --rule "seq[1],action[permit]" --set-
largecommunitylist-delete lgcomm1
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| rmap1          | 1       | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results

efa policy route-map-set create --name rmap1 --rule "seq[1],action[permit]" --set-
largecommunity 50:50:50,additive
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| rmap1          | 1       | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

2. Verify the following configuration on SLX devices:

```
SLX# show running-config route-map
route-map rmap1 permit 1
 set large-community 50:50:50 additive
 set large-community-list lgcomm1 delete
!
```

3. Run the following command to delete a route map set:

```
efa policy route-map-set delete
```

Example:

```
efa policy route-map-set delete --name rmap1 --rule "seq[1],action[permit]" --set-
largecommunity 50:50:50,additive --set-largecommunitylist-delete lg1
+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| rmap1          | 1       | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.159 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

Drift and Reconcile (DRC), Idempotency for Route Map Configuration for Large Community List

The following table describes the various attributes of route map for which DRC and idempotency is supported. A drift is identified if any of the fields below are modified by user through SLX CLI or other management tools. Reconcile operation pushes the intended configuration to SLX, therefore, bringing the SLX configuration in sync with XCO.

If you create an entry for idempotency which already exists in XCO, the system shows an error message stating that entry already exists.

Field	Identify Drift	Reconcile config	Idempotency	Comments
Update large community-list name in match criteria	Yes	Yes	No	Reconcile the large community list name
Large community list match criteria deleted	Yes	Yes	NA	Reconcile the match criteria for large community list
Update large community-list name in set criteria	Yes	Yes	No	Reconcile the large community list name
Large community list set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for large community list
Update large community attribute in set criteria	Yes	Yes	No	Reconcile the large community attribute
Large community attribute set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for large community attribute

Field	Identify Drift	Reconcile config	Idempotency	Comments
A different match criteria NOT supported by XCO is added through OOB	No	No	NA	
A set criteria NOT supported by XCO is added through OOB	No	No	NA	
Route map is created through OOB and this is not present or created by XCO	No	No	NA	

Policy Configuration Rollback

Perform the configuration rollback on route map, community list, and extcommunity list.

Policy Incremental Updates

The first type of configuration change is incremental updates to already provisioned objects, such as adding or removing rules or augmenting the contents within the rule (adding matches or sets). Ensure that the incremental update configuration is successful on all associated devices or not installed at all (rollback). In this scenario, following are the configuration output:

```
efa policy route-map-set create --name foo --rule seq[10],action[permit] --set-community
6550:125,internet,local-as

+----------------+---------+--------+
| Route Map Name | Seq num | Action |
+----------------+---------+--------+
| foo            | 10      | permit |
+----------------+---------+--------+
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.161 | Failed  |Some Err|                 |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Rollback|        |                 |
+---------------+---------+--------+-----------------+
Device Results
--- Time Elapsed: 10.33886575s ---
```

In the above output, the configuration failed on .161 but was successful on .162. However, since the policy change was unsuccessful on .161 the configuration is rolled back on .162. The Result of “rollback” indicates that the configuration was or is compatible with the configured device. It is possible that during the “Rollback” operation the configuration or Unprovisioning action fails. In this scenario the Result will also be designated as “fail”

but the cause of the failure, ie err message, will be contained ikn the “Rollback reason column”.

When performing “remove updates” on content within a policy the command is pre-validated to ensure none of the specified rules are not managed by XCO (XCO will not delete any policy information created Out of Band by the users). In this scenario, the CLI is errored out without proceeding to remove configuration from device or XCO DB. When the error is encountered, the user can either:

- Remove the device that contains the OOB configuration (see Policy Device Membership Updates on page 433)
- Remove the configuration from the device by using its native CLI.

Policy Device Membership Updates

The second type of policy configuration is the association or disassociation (add-device or remove-device) of a policy object to a device or list of devices.

Add Device

Add device operation is similar to the policy incremental updates. Ensure that the addition of the policy is successful on all specified device or not installed on any devices. The transaction status is reported identically as described in the Policy Incremental Updates on page 432. The following output shows that the addition of .162 would have been “successful” but was rolled back due to the failure of .161.

```
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.161 | Fail    |Some Err|                 |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Rollback|        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

Remove Device

Remove device is considered “best effort”. If any one of the specified devices “fails” the transaction there is no action taken and the configuration is still removed from all other devices whether successful or not. Unlike increment rule deletes there is no pre-validation or restriction with respect to devices that contain OOB create rules. If a device contains OOB entries they are ignored by XCO and no attempt to delete the configuration form the device will be made. For device remove XCO only removes configuration from the device and it’s internal DB for objects that XCO created.

The output of device addition and removal is identical to the output described in the Policy Incremental Updates on page 432. The following table depicts that the command was run for two devices to be removed. The device .161 failed for “Some Err”, the device may or may not still contain the configuration. However, the DB will still contain a mapping between the failed device and the specified policy object. Device

.162 was successfully removed and the DB and Device no longer contains the original configuration.

```
Route Map details
+---------------+---------+--------+-----------------+
|  IP Address   | Result  | Reason | Rollback reason |
+---------------+---------+--------+-----------------+
| 10.139.44.161 | Fail    |Some Err|                 |
+---------------+---------+--------+-----------------+
| 10.139.44.162 | Success |        |                 |
+---------------+---------+--------+-----------------+
Device Results
```

Provisioning Dependencies

There are certain system operation that will be treated as and “error” when configuration is executed the results in the configuration being rollbacked. The following is a list of system operations or states that will result in a rollback of a configuration request.

- DRC – If one of the specified devices within the command is actively performing DRC, the configuration will fail resulting in a rollback of configuration of all the specified devices.
- Admin Down – If one or more specified devices are in the administrative “Admin Down” state, the configuration will fail resulting in a rollback of configuration of all the specified devices.
- Firmware Download – If one of the specified devices within the command is actively performing a firmware download, the configuration will fail resulting in a rollback of configuration of all the specified devices.
- Concurrent configuration from a previous request – If one of the specified devices within the command is still actively performing a previous configuration request, all the subsequent request will fail resulting in a rollback of configuration of all the specified devices until the initial configuration is completed.

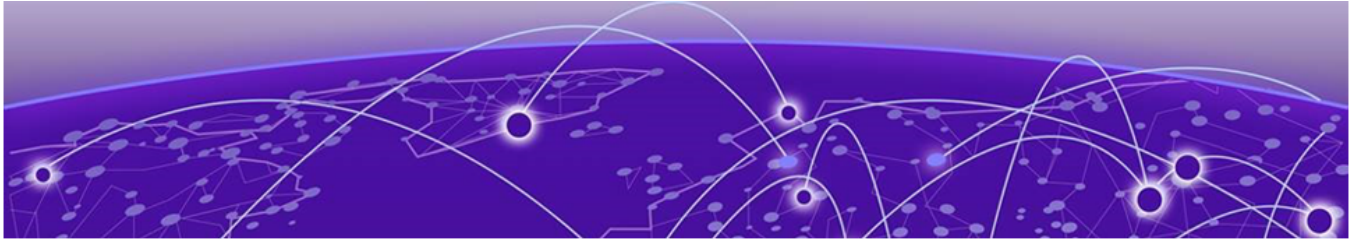

XCO Device Management

Device Image Management

XCO supports the following firmware download features.

- Firmware download with maintenance mode supporting the following:
  - Asynchronously launched operations
  - Sanity and pre-install script verification
  - Set convergence timeout, enable, and disable
  - Persisting the running configuration so that running configuration and maintenance mode configuration are preserved after reboot
  - Firmware download with the `no commit` option to enable restoration of firmware to a previous version
- Firmware host registration, with support for register, update, delete, and list operations
- Firmware download preparation, with support for add, remove, and list operations
- Firmware download with the `show` option, to display a table of devices in the fabric and their corresponding status

Limitations

- The device firmware must be SLX-OS 20.1.1 or later to support firmware download with maintenance mode for a hitless firmware upgrade.
- This feature assumes an existing host that contains SLX-OS firmware images ready to be downloaded.
- You can use this feature on a device where XCO TPVM is deployed, as long as you follow the instructions in *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.
- If you downgrade software from version 20.1.2a to 20.1.1, you must manually remove certificates.

Supported devices

The SLX-OS firmware download with maintenance mode is supported on the following SLX devices running SLX-OS 20.1.1 and later.

- SLX 9540
- SLX 9640
- SLX 9150-48Y
- SLX 9150-48XT
- SLX 9250
- SLX 9740

Hitless Firmware Upgrade

A hitless firmware upgrade uses the maintenance mode feature of the SLX device to gracefully divert traffic away from the device to alternate paths. The device can be put into maintenance mode and a firmware upgrade can be performed. The device can safely be rebooted and the new firmware activated without traffic loss. When the device is taken out of maintenance mode, traffic is allowed on the newly upgraded device.

Super-Spine Firmware Upgrade in Clos

1. The firmware on the first super-spine is downloaded.
2. Enabling maintenance mode on a super-spine involves the Border Gateway Protocol (BGP). The `graceful_shutdown` parameter is sent to all the super-spine's underlay neighbors (all connected spines). Each neighbor processes the `graceful_shutdown` and refreshes their routes to use the alternate path. Maintenance mode is enabled on the first super-spine and traffic is diverted to the second super-spine.
3. The `running-configuration` is saved on the first super-spine to preserve all current configurations including the maintenance mode enable configuration.
4. The device is rebooted for firmware activation without traffic loss.
5. Once the new firmware is activated, maintenance mode can be disabled. The `graceful_shutdown` parameter is removed from all the underlay neighbors and traffic to the first super-spine is allowed again.
6. The `running-config` is persisted again to ensure the maintenance mode disabled state is retained.

   The same process can be carried out on the second super-spine to upgrade the firmware without traffic loss.

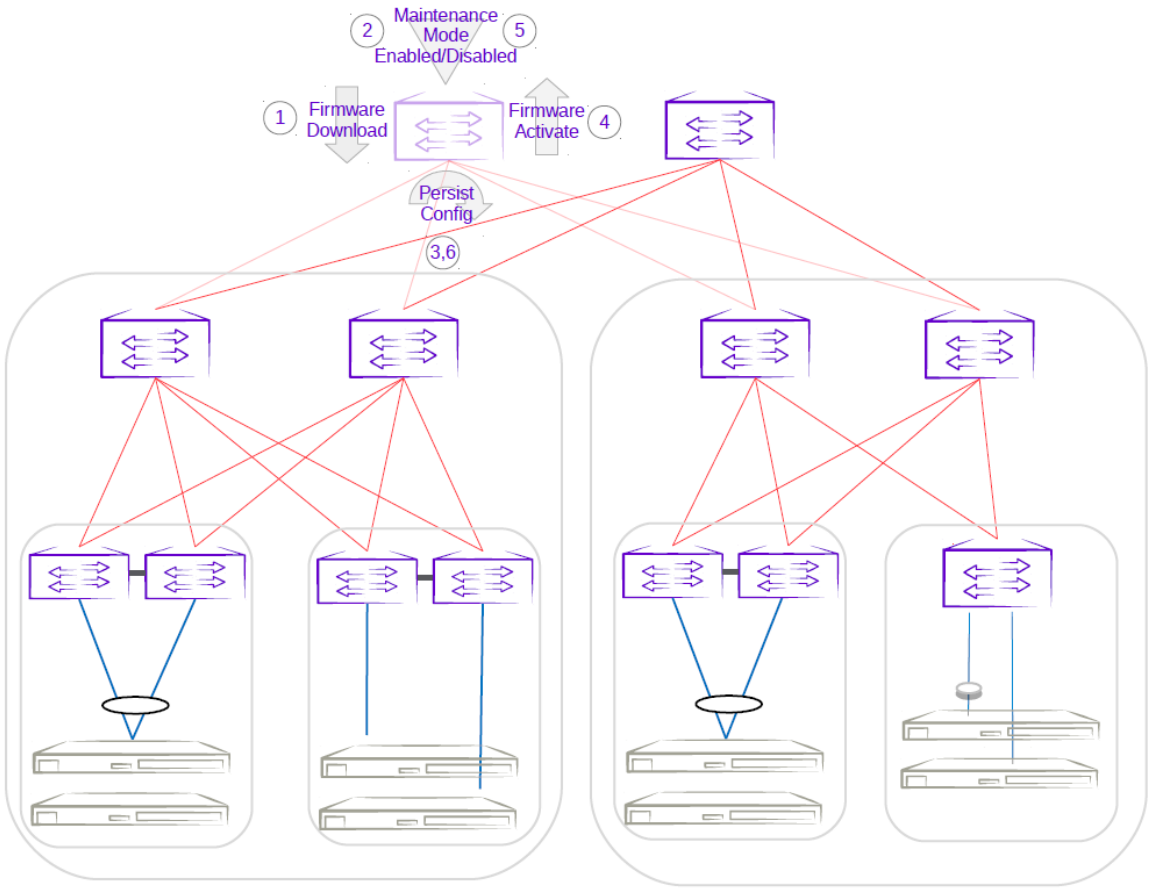

Figure 32: First super-spine firmware upgrade with maintenance mode

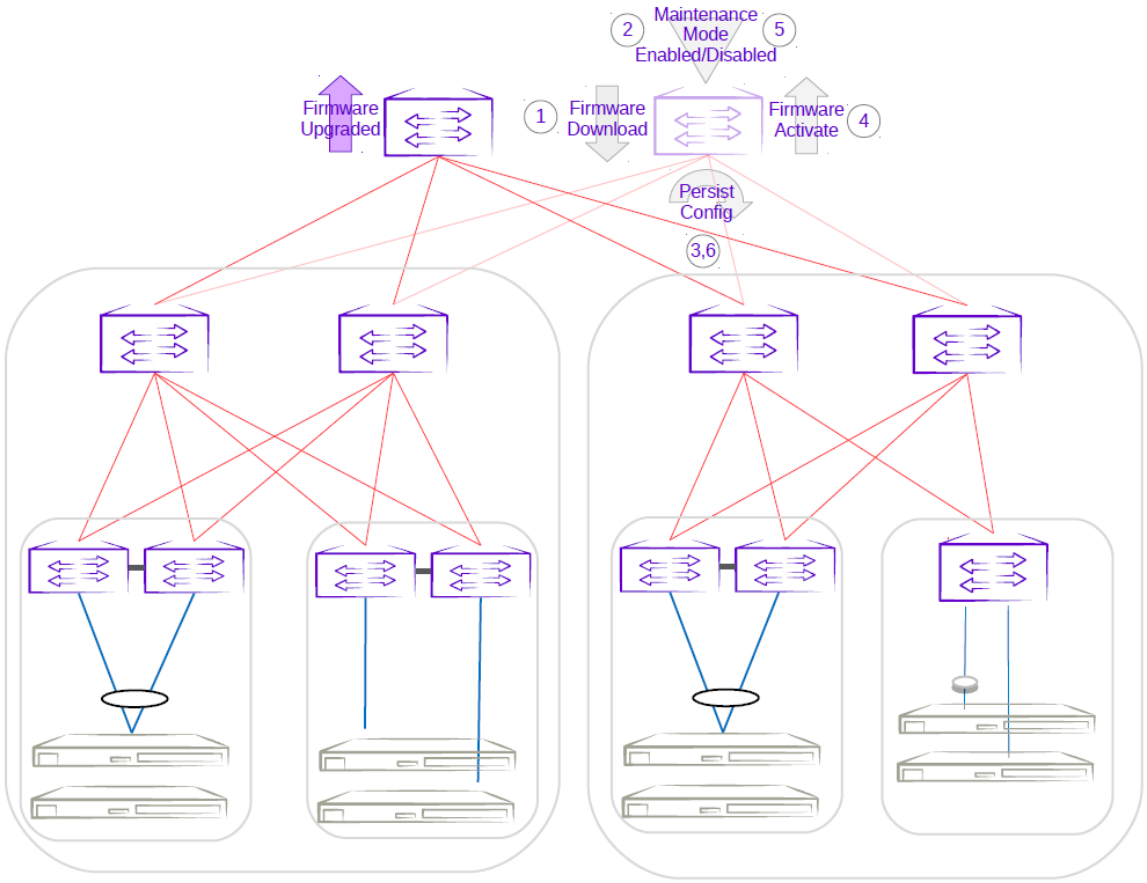

Figure 33: Second super-spine firmware upgrade with maintenance mode

Spine Firmware Upgrade in Clos

1. The firmware on the first spine is downloaded.
2. Enabling maintenance mode on a spine also involves the Border Gateway Protocol (BGP). The `graceful_shutdown` parameter is sent to all the spine's underlay neighbors (all leafs in the pod and super-spines). The neighbors no longer send traffic to the first spine going into maintenance mode and redirect traffic to an alternate path.
3. The `running-configuration` is saved on the first spine to preserve all current configurations including the maintenance mode enable configuration.
4. The device is rebooted for firmware activation without traffic loss.
5. Once the new firmware is activated, maintenance mode is disabled to allow traffic again through the upgraded spine.
6. The `running-config` is saved again to ensure the maintenance mode config remains disabled.

   The same process can be carried out on the second spine to upgrade the firmware without traffic loss.

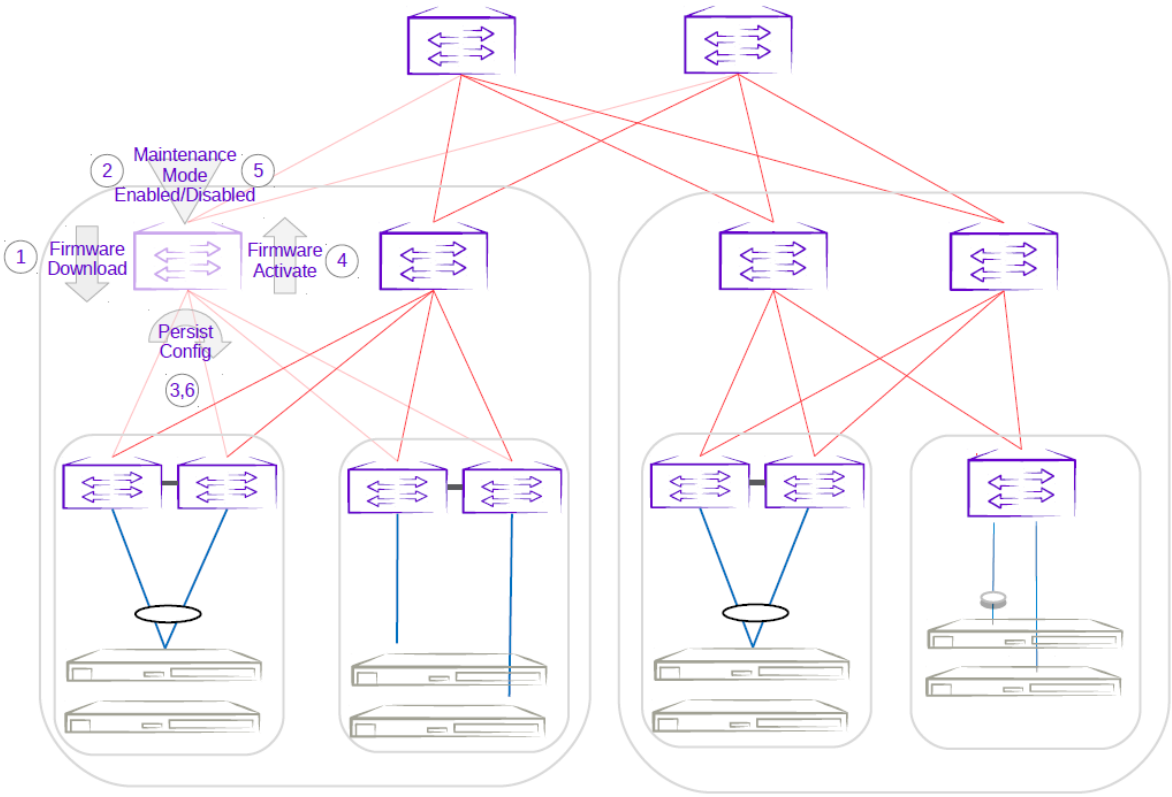

Figure 34: First spine firmware upgrade with maintenance mode

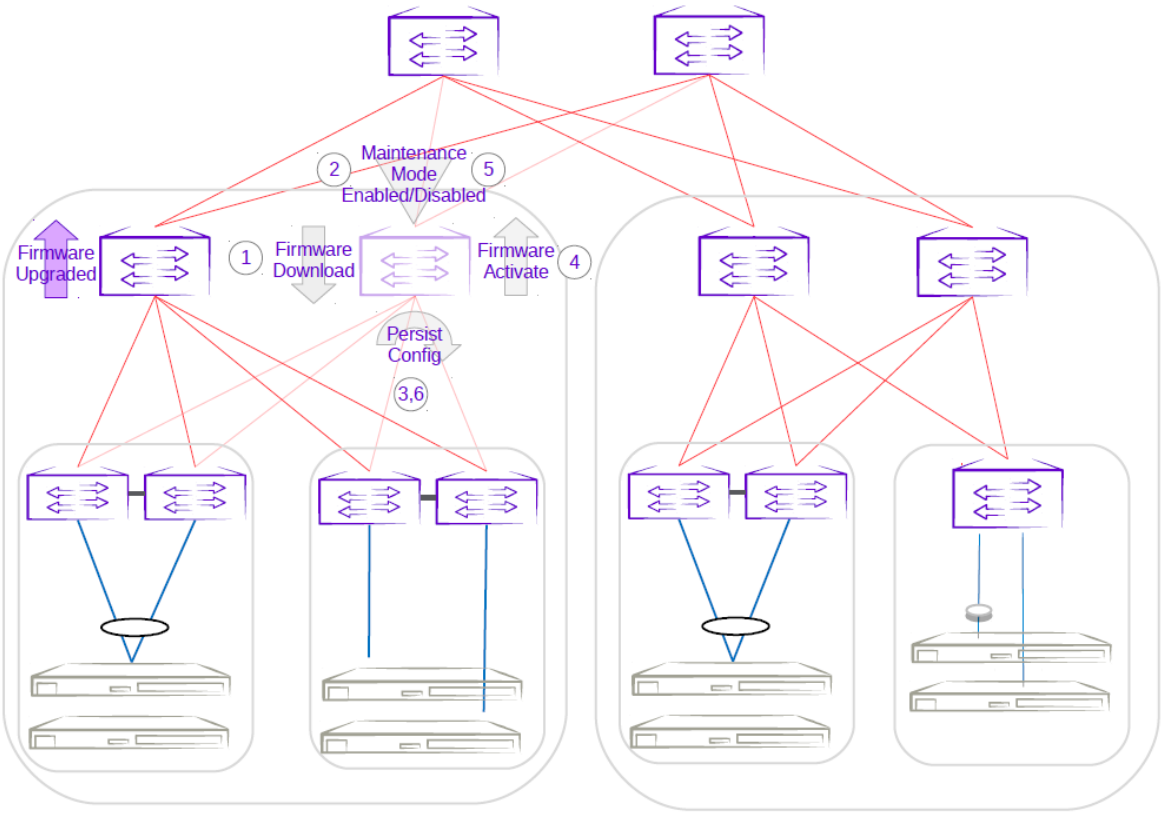

Figure 35: Second spine firmware upgrade with maintenance mode

Firmware Upgrade of an MCT Leaf Pair with Dual-Homed Servers in Clos

1. The firmware on the MCT leaf is downloaded.
2. Enabling maintenance mode on an MCT leaf involves the Border Gateway Protocol (BGP) and MCT or NSM. The `graceful_shutdown` parameter is sent to all the leaf's underlay neighbors (all spines in the pod). The neighbors no longer send traffic to the MCT leaf going into maintenance mode and redirect traffic from spines to the peer MCT leaf. MCT instructs the peer leaf to become the designated forwarder, ICL is shut down, and CCE ports for clients are also shut down. Traffic from dual-homed servers is redirected to the peer leaf. With maintenance mode enabled, traffic is completely redirected to the peer leaf.
3. The `running-configuration` is saved on the first MCT leaf to preserve all current configurations including the maintenance mode enable configuration.
4. The device is rebooted for firmware activation without traffic loss.
5. After the firmware is upgraded, the maintenance mode is disabled to allow traffic again through the upgraded MCT leaf.
6. The `running-config` is saved again to ensure the maintenance mode config remains disabled.

   The same process can be carried out on the second MCT leaf to upgrade the firmware without traffic loss.

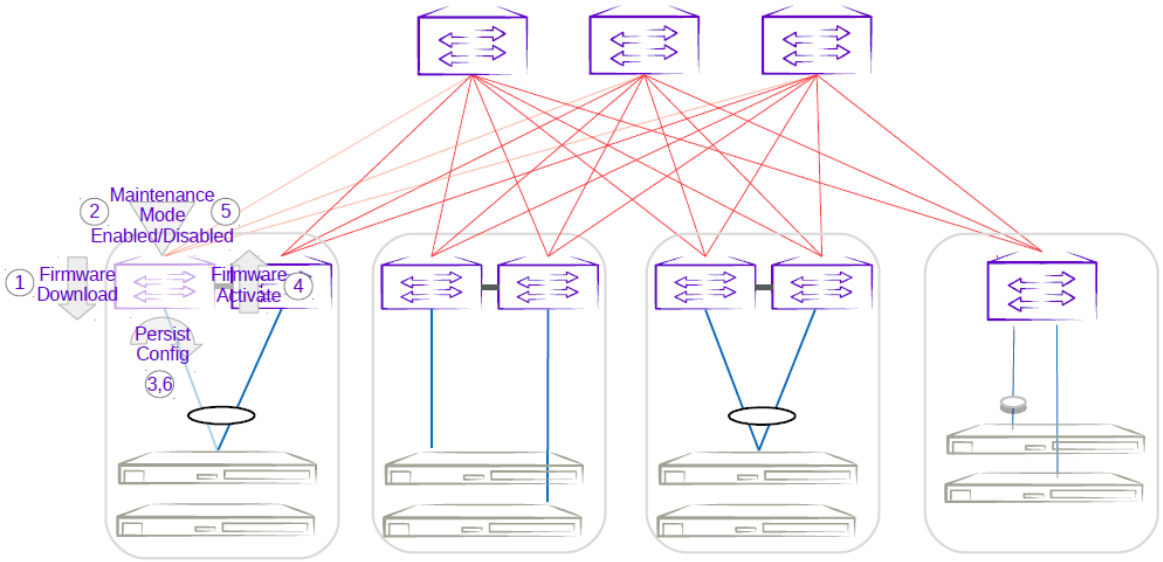

Figure 36: First MCT leaf firmware upgrade with maintenance mode

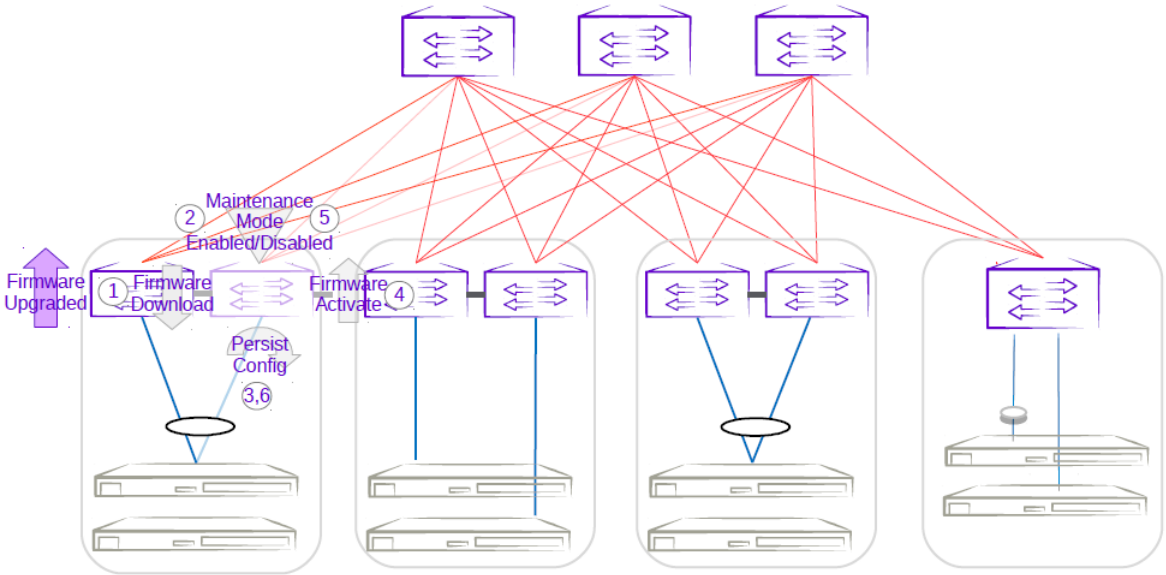

Figure 37: Second MCT leaf firmware upgrade with maintenance mode

Firmware Upgrade of a Three-Rack Centralized MCT Pair in Small Data Center

1. The firmware on the MCT leaf is downloaded.
2. Enabling maintenance mode on one of the leafs in the centralized MCT leaf pair follows the same behavior as the MCT leaf pair in a Clos topology. The only difference is the iBGP Layer 3 backup link between MCT leaf pairs. Maintenance mode results in the traffic being redirected to the peer leaf in the centralized MCT leaf pairs.
3. The `running-configuration` is saved on the first MCT leaf to preserve all current configurations including the maintenance mode enable configuration.

4. The device is rebooted for firmware activation without traffic loss.
5. After the firmware is upgraded, the maintenance mode is disabled to allow traffic again through the upgraded MCT leaf.
6. The `running-config` is saved again to ensure the maintenance mode config remains disabled.

   The same process can be carried out on the second MCT leaf to upgrade the firmware without traffic loss.

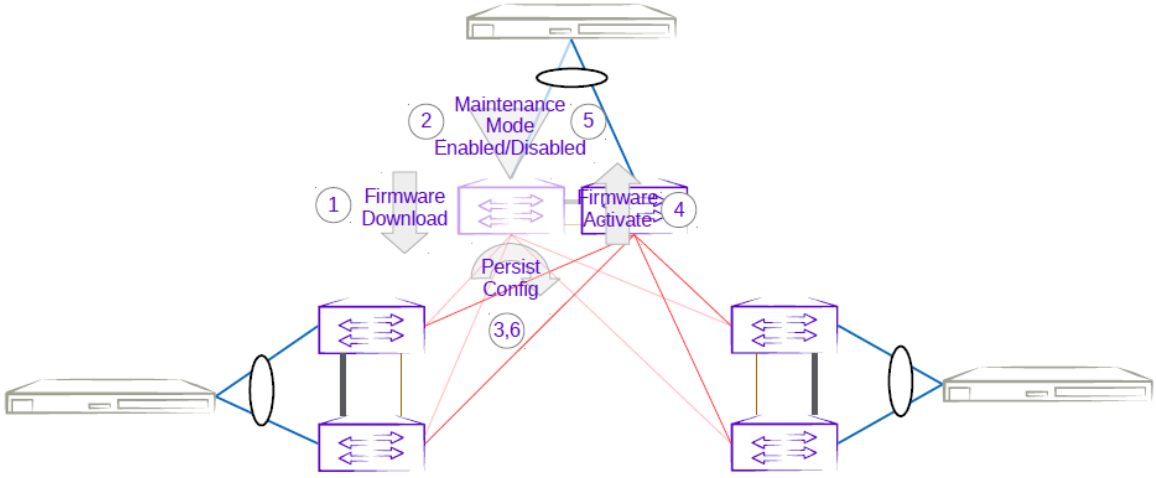

Figure 38: Three-rack centralized first MCT leaf firmware upgrade with maintenance mode

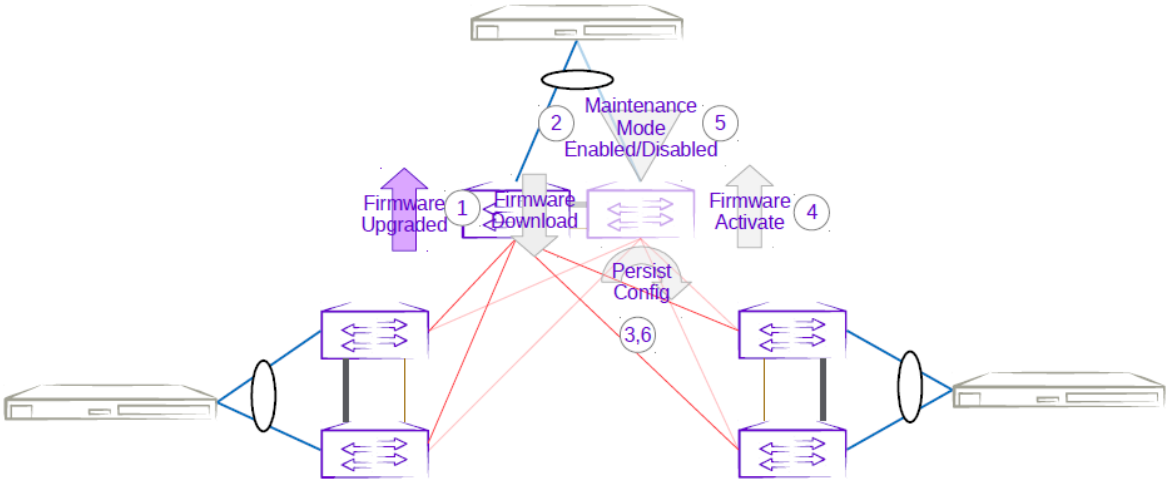

Figure 39: Three-rack centralized second MCT leaf firmware upgrade with maintenance mode

Firmware Upgrade of a Three-Rack Ring MCT Pair in Small Data Center

1. The firmware on the MCT leaf is downloaded.
2. Enabling maintenance mode on one of the leafs in a three-rack ring MCT leaf pair follows the same behavior as the MCT leaf pair in a Clos topology. The only difference

is the iBGP Layer 3 backup link between MCT leaf pairs. Maintenance mode results in the traffic being redirected to the peer MCT leaf.

3. The `running-configuration` is saved on the first MCT leaf to preserve all current configurations including the maintenance mode enable configuration.
4. The device is rebooted for firmware activation without traffic loss.
5. After the firmware is upgraded, the maintenance mode is disabled to allow traffic again through the upgraded MCT leaf.
6. The `running-config` is saved again to ensure the maintenance mode config remains disabled.

   The same process can be carried out on the second MCT leaf to upgrade the firmware without traffic loss.

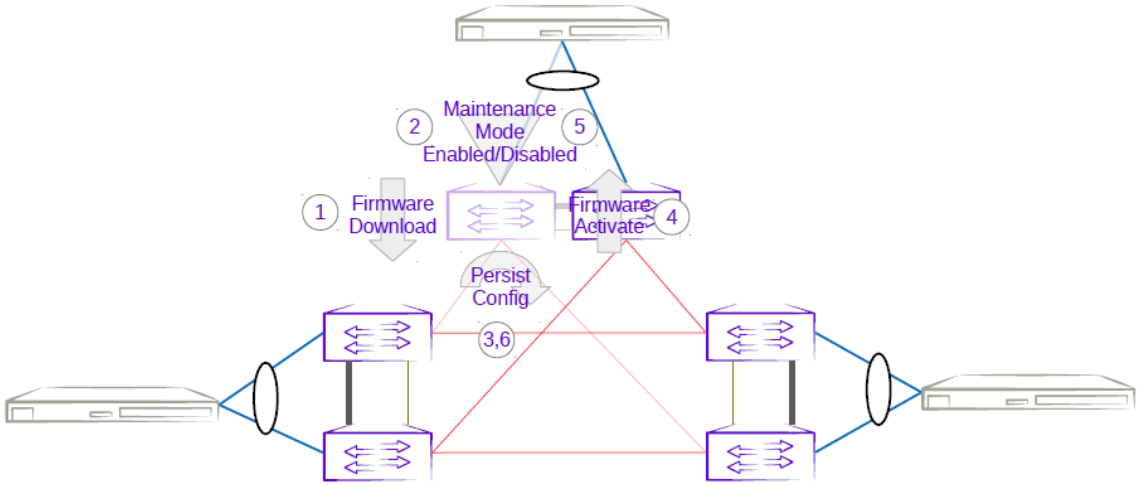

Figure 40: Three-rack ring first MCT leaf firmware upgrade with maintenance mode

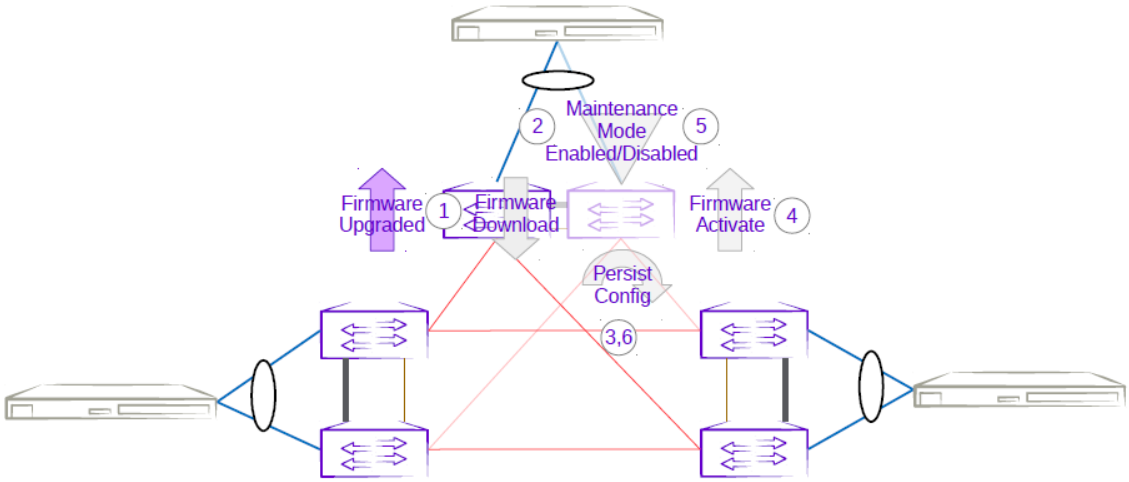

Figure 41: Three-rack ring second MCT leaf firmware upgrade with maintenance mode

Firmware Download

Use this topic to complete the firmware download and upgrade on fabric devices.

For more information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Firmware Upgrade with Minimal Traffic Loss

When updating firmware on a device, you typically begin by putting the device into maintenance mode so that traffic is diverted away from the device onto alternate paths. After performing the update, reboot the device, activate the new firmware, and take the device out of maintenance mode.

Alternatively, if it is not necessary to divert traffic away from the device, you can leave the device in active mode while updating the firmware. This enables the firmware to download faster.

Firmware Download Restart on HA Failover or Inventory Service Restart

Starting with EFA version 2.5.5, an in-progress firmware download restarts automatically if a high availability (HA) failover occurs, or if the inventory service restarts. This simplifies the process of preparing for the firmware download.

Firmware Download Implicit Fullinstall Support

The firmware download process automatically detects which devices require a firmware download fullinstall. You are warned that fullinstall is going to begin, because the fullinstall takes more time to complete. You do not need to provide any extra input outside of the normal prepare command.

XCO firmware download implicitly uses the “no reboot” option when running the firmware download command on the device. The SLX firmware download fullinstall command supports “no reboot” option starting with SLXOS 20.2.3ea. This is the minimum required SLX version which should already be installed on the device for XCO firmware download to perform fullinstalls to a later SLX version on the device.

SLX firmware download does not support nocommit and fullinstall options specified together, so XCO firmware-download reports an error if there are devices requiring a fullinstall and the `--noAutoCommit` flag has been specified.

1. The **`firmware download prepare`** command shows a warning when XCO detects that the device requires a full installation. The prepare is still successful.

   Example: [Supposing `firmware download prepare` detects a fullinstall is required for 10.20.246.4]

```
efa inventory device firmware-download prepare add --ip 10.20.246.4 --firmware-host
10.20.241.101 --firmware-directory /team/ztp/build/slx/slxos20.1.2e

+-------+------+-----+--------+---+----+-----------------+---------+---------------
+--------+------------------+
```

```
|IP     | Host |Model|Chassis |ASN|Role|Current Firmware  |Firmware |    Firmware
|Target  | Last Update Time  |
|Address| Name |     |Name    |   |    |                  |Host     |    Directory
|Firmware|                  |
+-------+------+-----+--------+---+----+------------------+---------+----------------
+--------+--------------------+
|10.20  | NH-  |3012 |SLX9250 |0  |Leaf|20.2.3slxos20.2.3b|10.20.241|/team/ztp/build/|
20.1.2e |2021-03-30 22:12:52|
|.246.4 | Leaf2|     |-32C    |0  |    |_210309_0732      |.241.101 |slx/
slxos20.1.2e|        |+0000 UTC          |
+-------+------+-----+--------+---+----+------------------+---------+----------------
+--------+--------------------+
Firmware Download Prepare Add Details
Validate Firmware Download Prepare Add [success]

Warning(s)
    10.20.246.4: Full installation is required to firmware download 20.1.2e from
existing version. Firmware download will take longer to complete on device 10.20.246.4
```

2. The **`firmware download execute`** command shows an error when `--noAutoCommit` is specified and one or more devices require a full installation.

   Example: [Supposing **`firmware download execute --noAutoCommit`** is issued and fullinstall is required for 10.20.246.4]

```
efa inventory device firmware-download execute --fabric non_clos --noAutoCommit
Firmware Download Execute [failed]
    10.20.246.4: Device 10.20.246.4 cannot perform firmware download with noAutoCommit
and fullinstall requirement
```

XCO Command Blocking during Firmware Download

Before starting the firmware download, XCO verifies that all system services are not currently busy. If the verification does not complete within 2 minutes and 30 seconds, XCO displays an error. Retry the firmware download later.

Inventory Command Blocking

The following inventory commands are blocked when a firmware download is in progress for a specific device:

- Network Essentials (Native CLIs)
- Device Execute CLI
- Device Delete

Fabric Command Blocking

- The following fabric commands are blocked when at least one of the device is in `fwdl-in-progress` state:
  - **Fabric device add**

    Example: [Supposing firmware download is in progress for 10.20.246.1]

```
efa fabric device add --ip 10.20.246.2 --rack rack1 --name non-clos --username
admin --password password

Error:  Device(s) 10.20.246.1 are going through firmware download
```

  - **Fabric device remove**

Example: [Supposing firmware download is in progress for 10.20.246.1]

```
efa fabric device remove --name non-clos --ip 10.20.246.1,10.20.246.2
         Remove Device(s) [Failed]
         Removal of  device with ip-address =  [Failed]

Reason: Device(s) 10.20.246.1 are going through firmware download
```

- **Fabric configure**

  Example: [Supposing firmware download is in progress for 10.20.246.1]

```
efa fabric configure --name non-clos

Error:  Device(s) 10.20.246.1 are going through firmware download
```

- The following fabric commands are allowed even when the devices are in `fwdl-in-progress` state:
  - Fabric device remove with “no-device-cleanup" flag
  - Fabric delete with “force” option
  - Fabric topology show physical/underlay/overlay to display the output by excluding the devices in `fwdl-in-progress-state`

Tenant Command Blocking

- If the target device list for a particular tenant operation has at least one device in `fwdl-in-progress` state, then the entire operation is rejected with an error to the user.

  Example: [Supposing firmware download is in progress for 10.24.80.158]

```
efa tenant po create --name po1 --tenant t1 --speed 10Gbps --negotiation active --port
10.24.80.158[0/3], 10.24.80.159[0/3]

PortChannel creation failed:
      Error: Firmware download is in progress for the devices [10.24.80.158]
```

- If the target device list for a particular tenant operation has no device in `fwdl-in-progress` state, then the operation proceeds as usual.

  Example: [Supposing firmware download is in progress for 10.24.80.158]

```
efa tenant po create --name po1 --tenant t1 --speed 10Gbps --negotiation active --port
10.24.80.159[0/3]

PortChannel creation succeeded
```

Failures During Group-based Firmware Download Execution

Failures can occur during the firmware download process such as network connectivity issues. Any such failures for a device during the group-based firmware download execution remain in the error state, and the execution proceeds to the next group

by default. The **`firmware download execute`** command contains a new `--group-execution` parameter to control this behavior. The two policies are:

- **continue-on-error (default)**: If any device results in an error during the firmware download process, firmware download execution will continue and process all the remaining groups. The overall status is failed.
- **stop-on-error**: If any device results in an error during the firmware download process, the firmware download execution will not proceed to the next group. Any devices in the remaining groups are left in a prepared state. The overall status is failed.

In the event of an error, you can restart the firmware download operation.. The execution automatically restarts on the firmware download process on failed device(s) per group

Group-based Firmware Download Restore

When the `--noAutoCommit` option is used with the firmware download execution, the device retains the previous firmware image. It enables you to test out the new firmware and decide to keep it by issuing a firmware download commit or go back to the previous firmware image by issuing a firmware download restore.

You can also apply the firmware download commit to all the devices in the fabric simultaneously. However, the firmware download restore implicitly reloads the device, so that the firmware download restore must be staged to prevent and minimize traffic loss.

You can invoke the firmware download restore for a fabric or for a set of IP addresses. In both the cases, the restore is applied in the same group order defined by the prepared list.

IP addresses based restore must be called for devices in the same fabric.

Group execution policy is inherited from what was specified in the firmware download execute CLI.

Fabric-wide Firmware Download

About This Task

Follow this procedure for upgrading the firmware of devices in a Clos fabric. It describes how to upgrade the device of standby XCO node and MCT leaf pairs, force a failover to change the active node to standby, and then upgrade the SLX of new standby node and remaining MCT leaf pairs.

To upgrade firmware in a small data center configuration, see *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

Procedure

1. Prepare and run the firmware download on the devices in the fabric, in batches. In batch-1, add the device that hosts the standby node and devices on right side of the

fabric. The diagram that follows illustrates the right and left devices in the batches of a fabric.

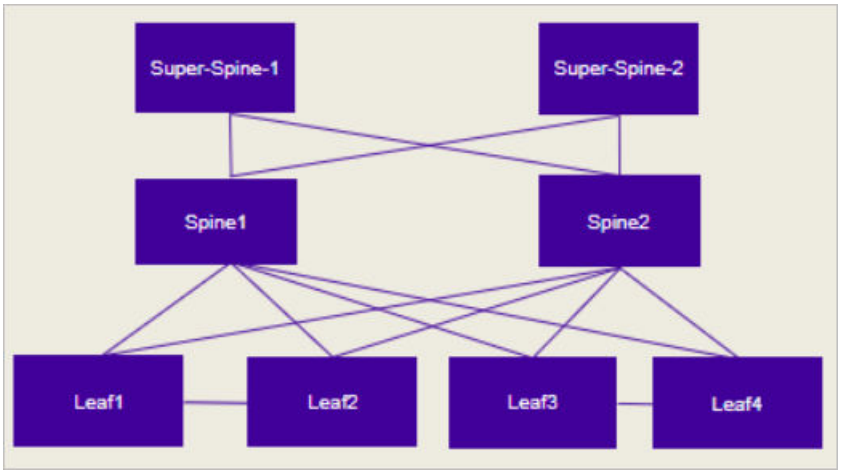

Figure 42: Batches for fabric-wide update

a. Prepare the firmware download.

```
$ efa inventory device firmware-download prepare add --fabric <fabric name> --
firmware-host <IP of firmware download host>
--firmware-directory <path to target firmware build>
```

The command returns the following information in a table: IP address, host name, model, chassis name, ASN, role, current firmware, firmware host, firmware directory, target firmware, and last update time.

b. Download the firmware with or without the `–noAutoCommit`, `–noMaintMode`, and `–drc` options, as desired.

```
$ efa inventory device firmware-download execute
--fabric <fabric name> --noAutoCommit –noMaintMode

Firmware Download Execute [success]

--noAutoCommit    Configure Auto commit in Firmware Download
--noMaintMode     Configure Maintenance Mode in Firmware Download
--drc             Configure a drift reconcilation operation is performed after
Firmware Download
```

c. Monitor the progress of the firmware download.

```
$ efa inventory device firmware-download show
--fabric <fabric name>
```

d. Repeat step c until the firmware download is complete.

Each time you repeat step c, the command returns a table that details the progress of the firmware download. The download is complete when the Update State column shows **Completed** and the Status column shows **Firmware Not Committed** when `–noAutoCommit` is used and **Firmware Committed** without `–noAutoCommit`.

2. Commit the firmware across all devices in the fabric.

```
$ efa inventory device firmware-download commit –fabric <fabric name>

OR

$ efa inventory device firmware-download commit –ip <IP address of all devices in
fabric>
```

The download is complete when the Update State column shows **Completed** on all devices and the Status column shows **Firmware Committed**.

Group-based Firmware Download Preparation

The firmware download prepare commands **`add`** and **`delete`** accept a new `--group <#>` parameter. Group creation and deletion is inferred by the existence of a prepared device in a group. These commands enable you to build and modify a custom prepared list for the entire fabric.

Fabric-based Firmware Download Preparation

The firmware download prepare **`add`** and **`delete`** command accept a new `--fabric <fabric name>` parameter. This parameter automatically generates a group-based prepared list for the entire fabric or delete any existing prepared list for the entire fabric.

You can review and further modify the auto-generated prepared list using the group-based prepare commands if required.

Clos Topology (3-Stage and 5-Stage)

The following rules for Clos topologies generate a fabric-based prepared list:

- The first group contains all non-MCT leaf devices and the MCT peer with the lower IP address from all MCT leaf and border-leaf devices.
- The second group contains the MCT peer with the higher IP address from all MCT leaf and border-leaf devices.
- The subsequent groups contains a single spine per group starting from the lowest to highest IP address of the spines.

The remaining groups contain a single super-spine per group starting from the lowest to highest IP address of the super-spines.

Small Data Center Topology (Centralized Rack and Rack Ring)

The following rules for small data center topologies generate a fabric-based prepared list:

- Due to potential loss of connectivity between racks (rack ring topology), all the devices are prepared for a serial upgrade (one device per group).
- The order is rack-by-rack, with the lower IP address peer followed by the higher IP address peer.

Group-based Firmware Download Execution

A single firmware download execution starts with the smallest group number performing the firmware download simultaneously on all devices prepared in the group and continue processing each group sequentially up through the largest group number. The group numbers need not be contiguously defined.

The prepared list remains until the entire group-based firmware download execution is completed and the firmware is committed or restored.

The following diagrams show an example of the group-based firmware download execution when prepared using the fabric-based auto-generated prepared list.

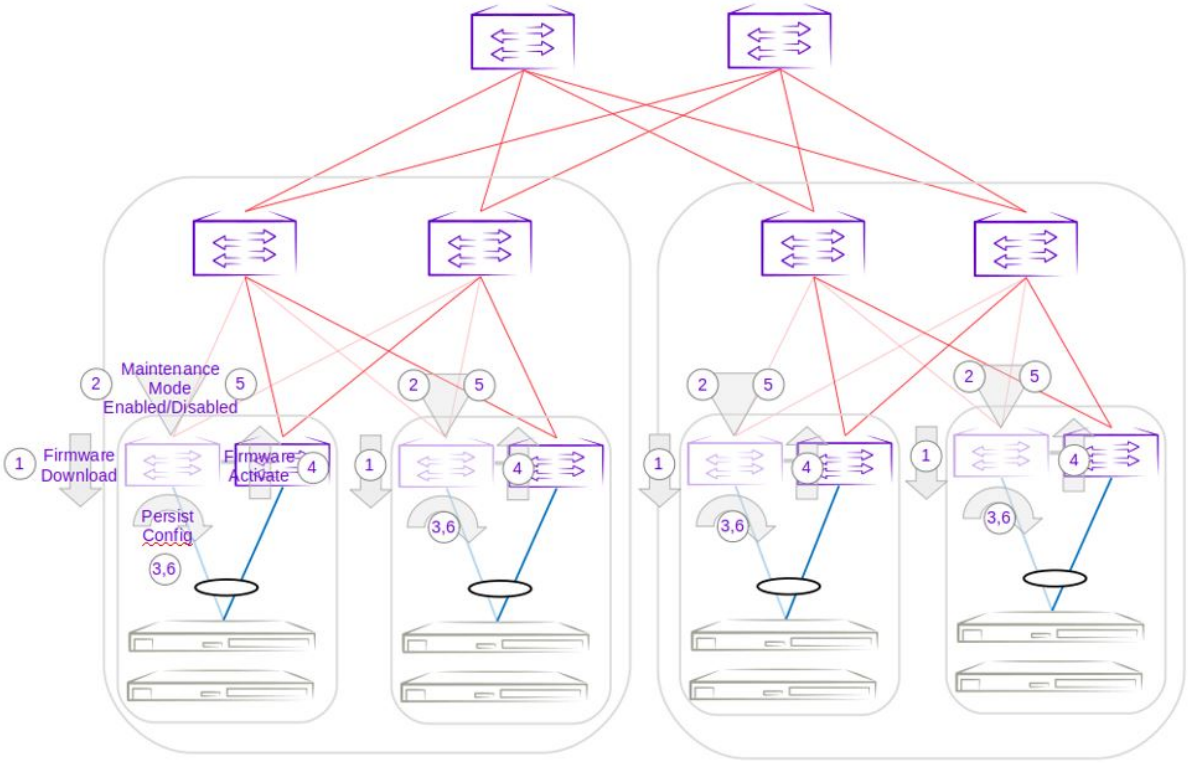

Figure 43: Group 1 - Firmware download execution of lower IP address MCT Peer leaf devices

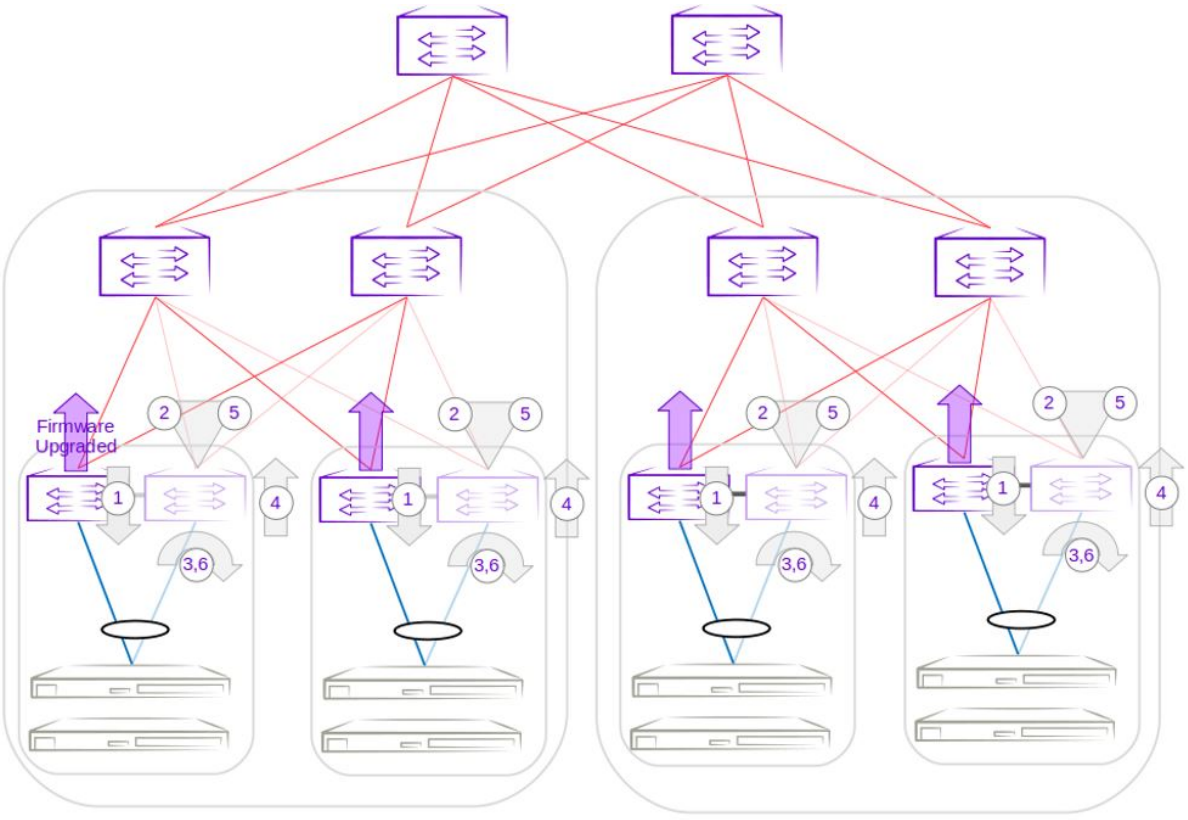

Figure 44: Group 2 - Firmware download execution of higher IP address MCT Peer leaf devices

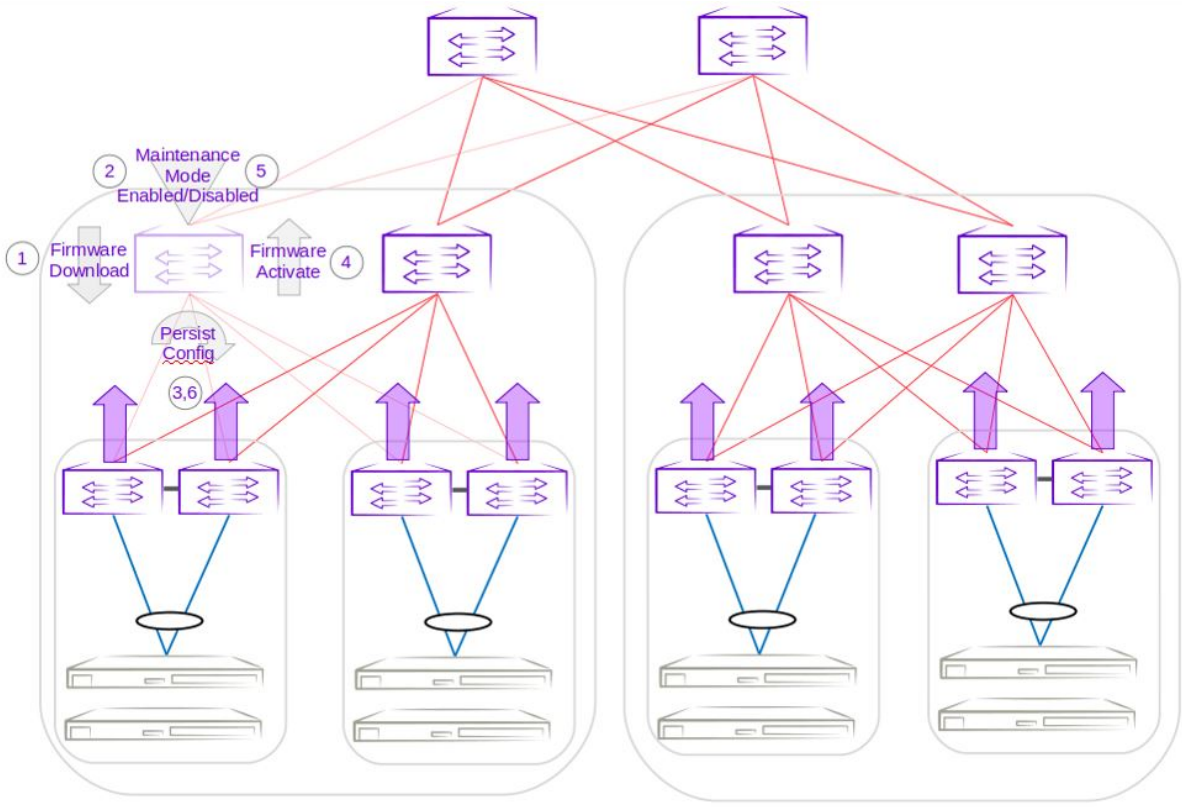

Figure 45: Group 3 - Firmware download execution of lowest IP address spine device

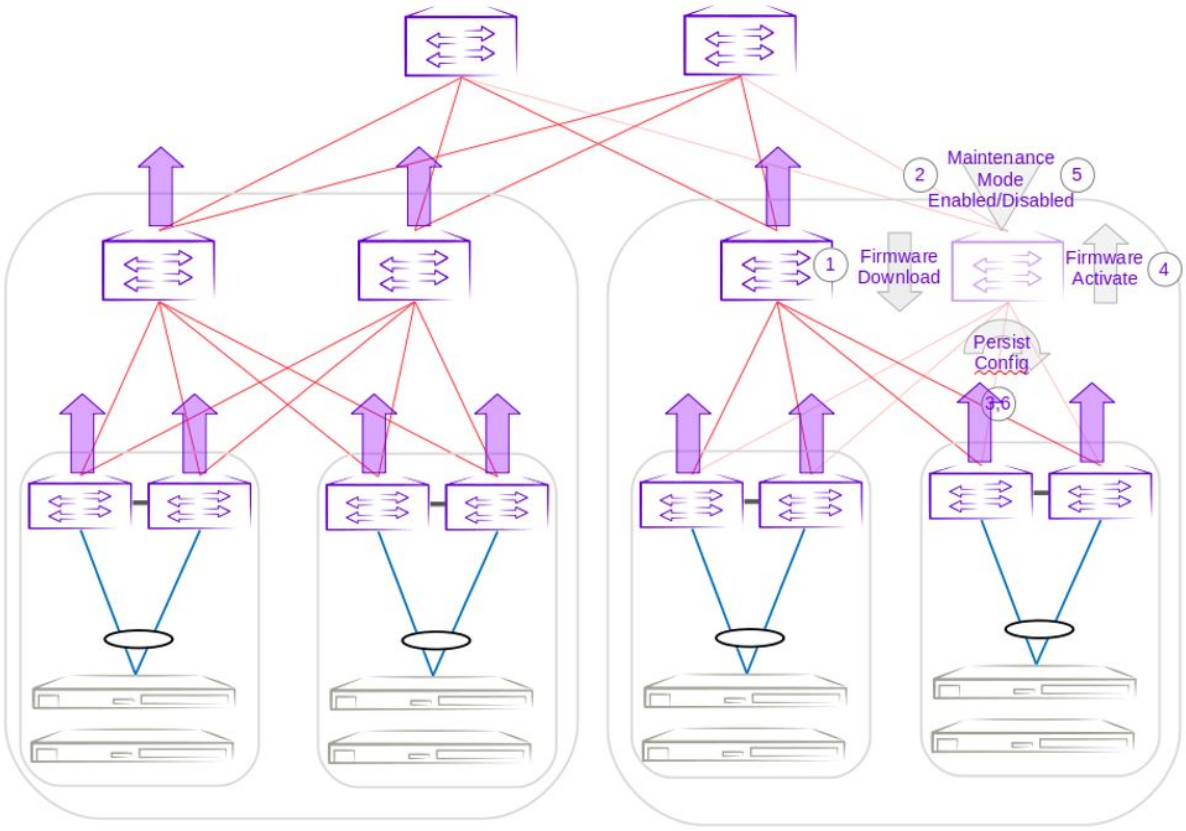

Figure 46: Group 4 - Firmware download execution of highest IP address spine device

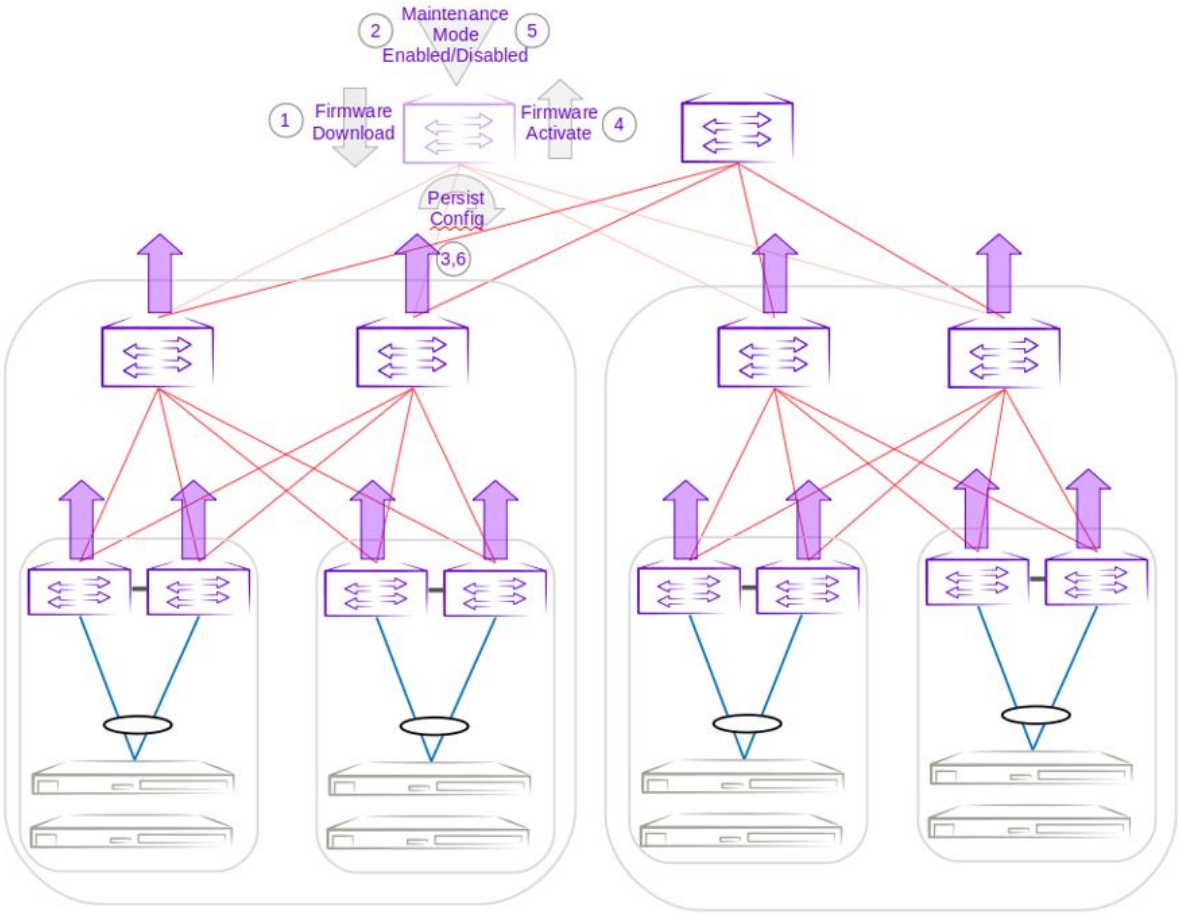

Figure 47: Group 5 - Firmware download execution of lowest IP address super-spine device

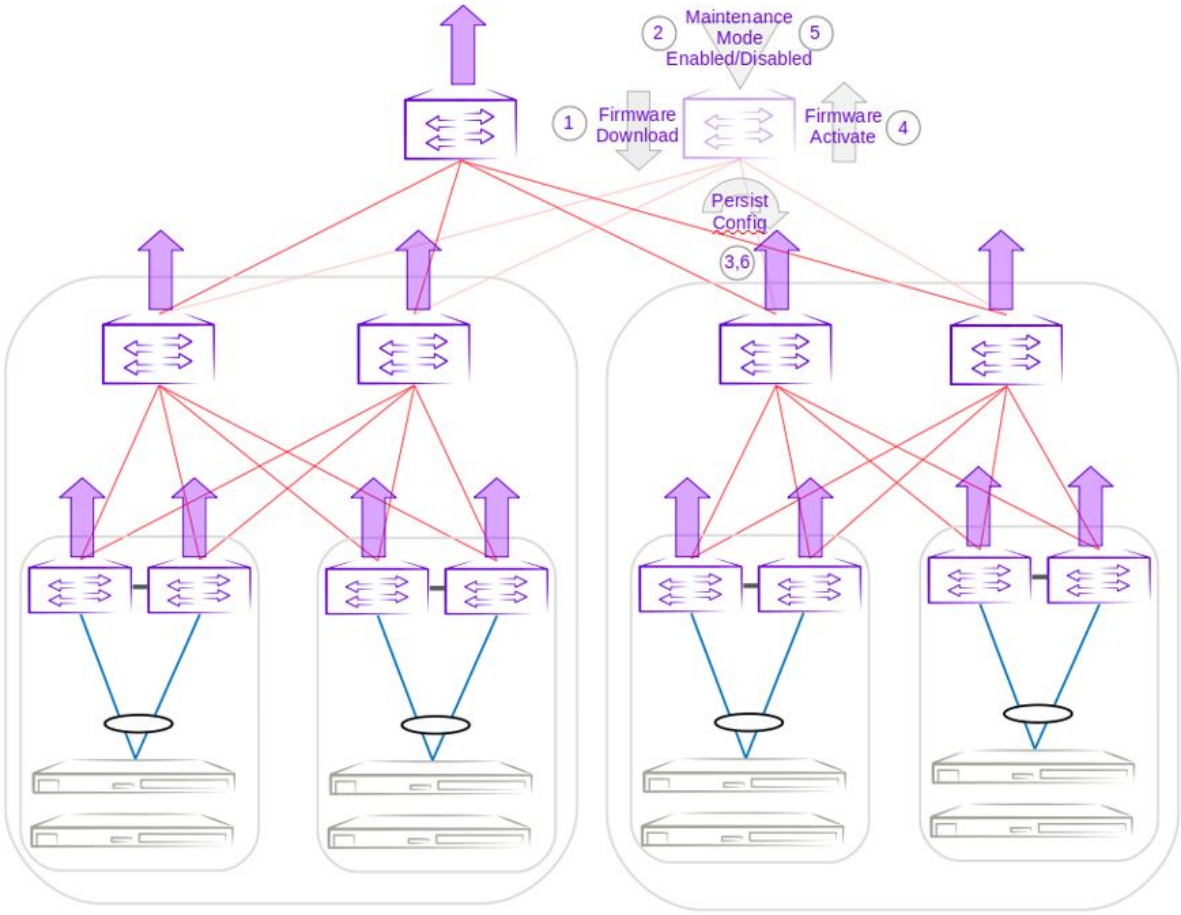

Figure 48: Group 6 - Firmware download execution of highest IP address super-spine device

Roll Back Device Firmware

You can roll back the firmware on the device when it is in "Firmware Not Committed" status.

About This Task

This is the recommended method for rolling back firmware when it is not committed. Run firmware restore on all devices in the fabric.

Procedure

Run the following command to restore the firmware across all devices in the fabric:

```
$ efa inventory device firmware-download restore –fabric <fabric name>
OR
$ efa inventory device firmware-download restore –fabric <IP address of all devices in fabric>
```

The download is complete when the Update State column shows **Completed** on all devices and the Status column shows **Firmware Committed**.

Traffic Loss Scenarios

Single Leaf

Traffic loss is expected when you upgrade a single leaf that is not in an MCT pair. Because there are no alternate paths for the single leaf, maintenance mode is not enabled. Only the configuration is persisted, and a firmware upgrade is carried out. A traffic loss warning is flagged when you upgrade a single non-MCT leaf.

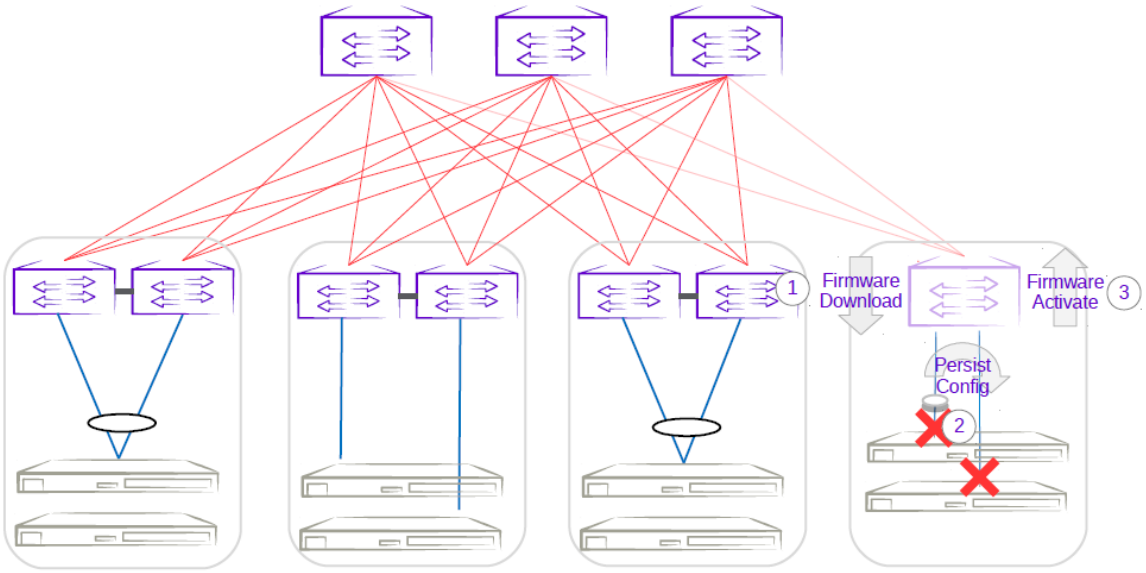

Figure 49: Single-leaf traffic loss

Single-Homed Server

Traffic loss is also expected for any singled-homed server. Detecting single-homed servers are not in the scope of this feature so a generic warning is provided at the start of a firmware download.

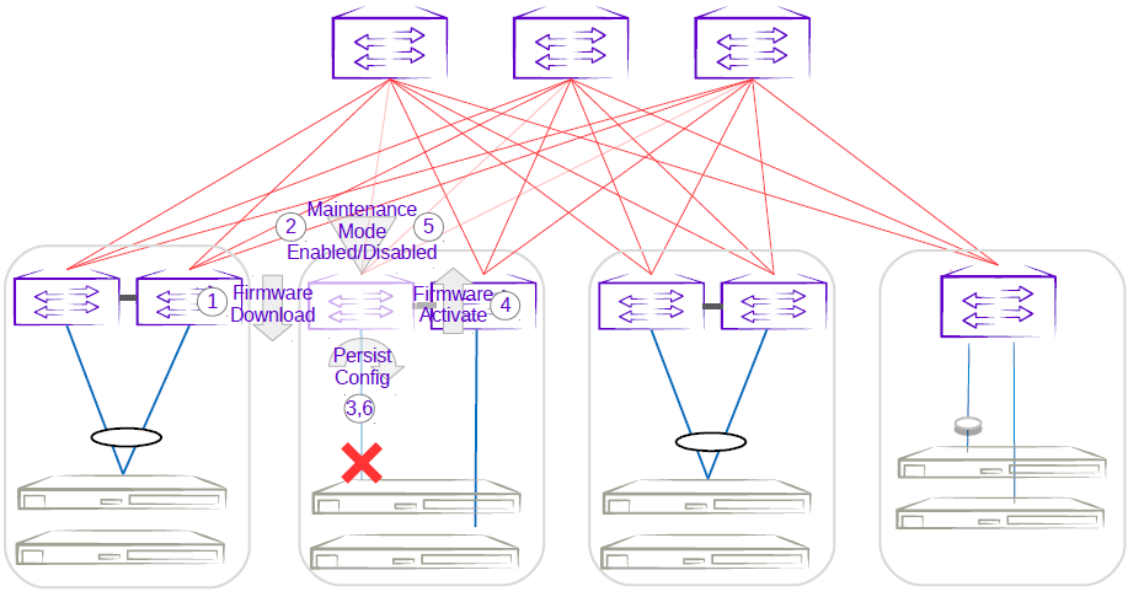

Figure 50: Single-homed server traffic loss

Non-Redundant Spine or Super-Spine

This is not a typical deployment, but traffic loss is expected in this scenario. Because no alternate paths exist for non-redundant devices, maintenance mode is not enabled for this case. A traffic loss warning is flagged when you upgrade non-redundant devices.

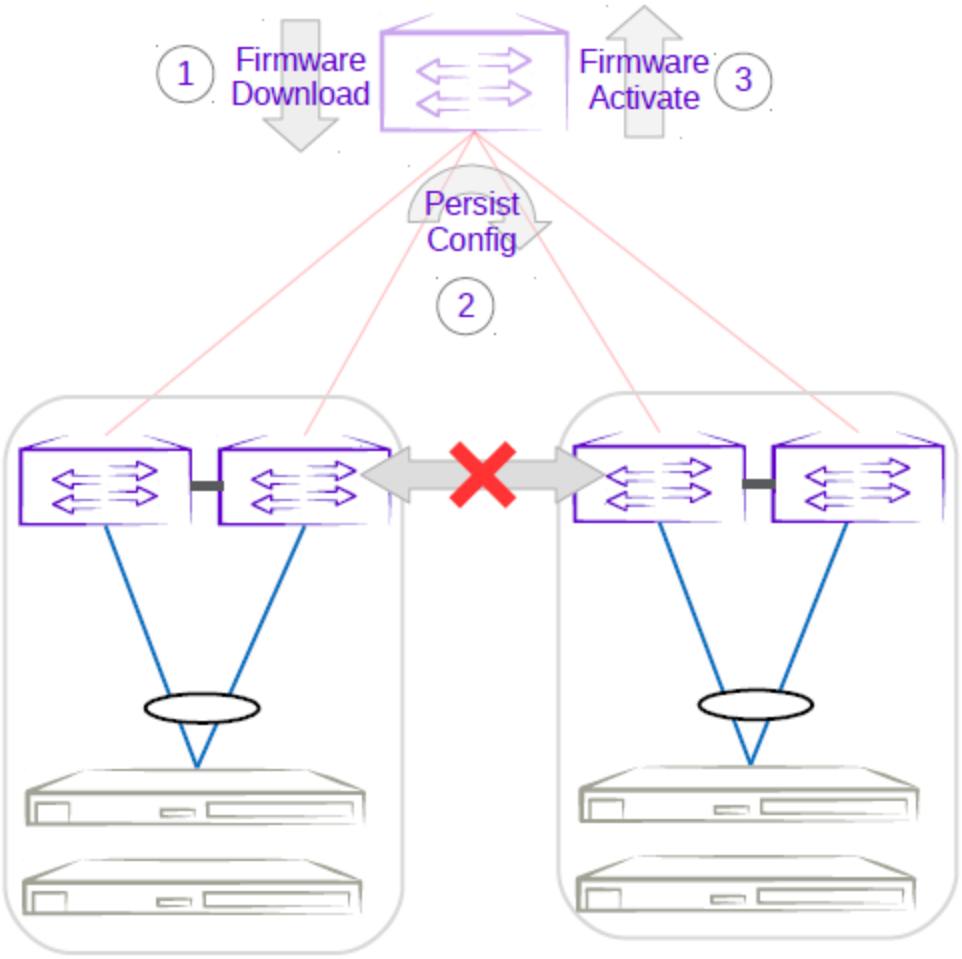

Figure 51: Non-redundant spine traffic loss

Device Health Management

Device Health Management (DHM) performs drift and reconciliation services, restoring fabric related configurations.

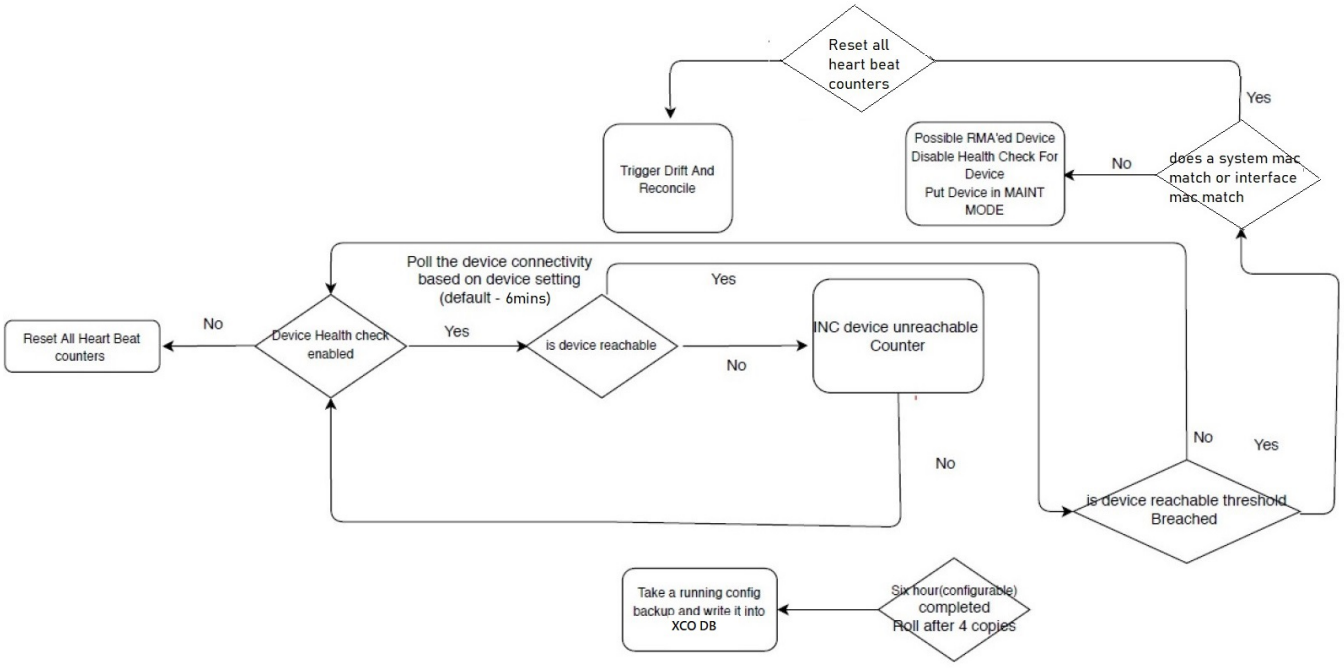

Figure 52: Device Health Management work flow

Monitor Device Health

The devices registered with XCO can be monitored for connectivity issues. If connectivity violates pre-defined thresholds, XCO starts drift and reconciliation.

About This Task

Follow this procedure to monitor device health.

Procedure

1. Enable device health check.

```
# efa inventory device setting update --ip 10.24.14.133 --health-check-enable yes
```

2. Configure health check interval.

```
# efa inventory device setting update --ip 10.24.14.133 --health-check-interval 30mins
```

3. Configure health check threshold.

```
# efa inventory device setting update --ip 10.24.14.133 --health-check-heartbeat-miss-threshold 2
```

4. View device health status.

```
# efa inventory device health status --ip 10.24.14.133
```

5. (Optional) Disable device health check.

```
# efa inventory device setting update --ip 10.24.14.133 --health-check-enable no
```

Device Configuration Backup and Replay

The Device Configuration Backup and Replay feature enables backup of the device configuration based on inventory device settings or user-run commands and REST APIs.

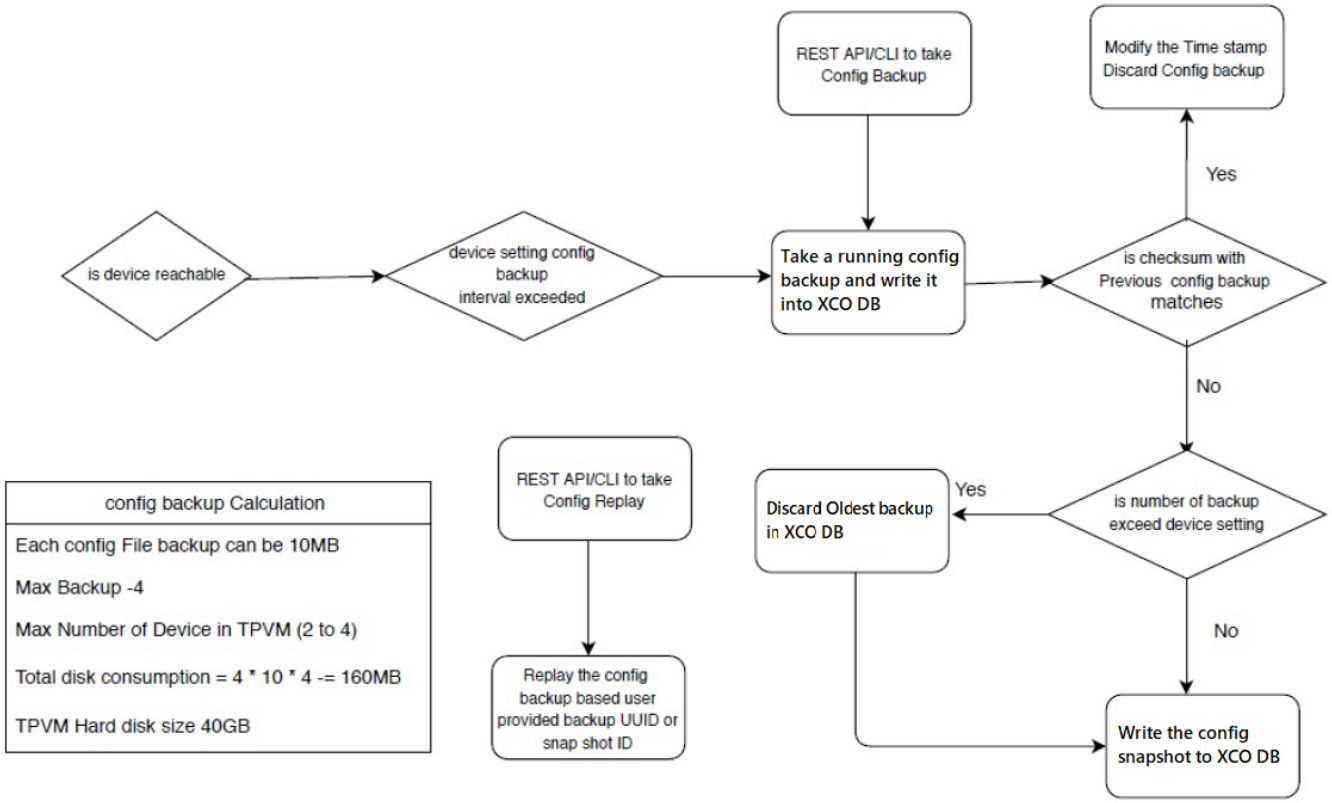

Figure 53: Workflow

Configure Backup and Replay

You can configure backup and replay of device.

About This Task

Use this procedure to configure backup and replay.

Procedure

1. Enable periodic config-backup.

```
# efa inventory device setting update --ip 10.24.14.133 --config-backup-periodic-
enable yes
```

2. Configure device backup.

```
efa inventory device setting update --ip 10.24.14.133 --config-backup-interval 30m
[3m-1800m, default 1440m]
# efa inventory device setting update --ip 10.24.14.133 --number-of-config-backups 2
[2-20, default 4]
# efa inventory config-backup execute --ip 10.24.14.133
```

3. View config-backup history.

```
# efa inventory config-backup history --ip 10.24.14.133
# efa inventory config-backup detail --uuid 1111-1111-1111 --show-config
# efa inventory config-backup detail --uuid 1111-1111-1111 --show-config --file-dump
<filename>
```

4. (Optional) Delete config-backup.

```
# efa inventory config-backup delete --key 10.24.14.133
# efa inventory config-backup delete --key 1111-1111-111
```

5. Determine the backup restore method:
   - To restore backup using the `startup-config` file, proceed to the next step.
   - To restore backup using the `running-config` file, go to Step 7.

   Note
   The `startup-config` backup restore method requires device reboot to restore the configuration.

6. Configure device replay using the appropriate command.
   - Config-replay without rebooting the device:

```
# efa inventory config-replay execute --ip 10.24.14.133 --uuid 1111-1111-111 --
startup-config --no-reboot
```

   - Config-replay with device reboot:

```
# efa inventory config-replay execute --ip 10.24.14.133 --uuid 1111-1111-111 --
startup-config
```

7. Configure device replay using `running-config`.

```
# efa inventory config-replay execute --ip 10.24.14.133 --uuid 1111-1111-111
```

8. View config-replay history.

```
# efa inventory config-replay history --ip 10.24.14.133
# efa inventory config-replay detail --uuid 1111-1111-1111
```

9. (Optional) Delete config-replay.

```
# efa inventory config-replay delete --key 10.24.14.133
# efa inventory config-replay delete --key 1111-1111-111
```

Return Material Authorization

With the Return Material Authorization (RMA) process, you can replace a faulty device with a new device that has the same configuration.

The high-level process is as follows. For specific steps and commands, see Replace a Faulty Device on page 462.

1. Verify prerequisites.
   - Periodic configuration backup must be enabled on all devices that may need RMA. This prerequisite ensures that you have the latest configuration file to be used for recovery.
   - Maintenance mode must be enabled upon reboot on all devices.
2. Remove the faulty device and replace it with the new device. Ports on this device must be administratively up and online.

   The ports on the new device must have the same connections as the old device. For example, if the old device Ethernet port1 and port2 went to port3 and port4 of another device, the new replacement device must have these exact same Ethernet port connections.

3. Configure the new device with the same management IP address and credentials as the old device.
4. Start the RMA process from the command line or with the REST API.

   Note
   As a best practice, run the **`efa inventory rma execute`** command with the configuration backup ID so that the configuration is properly restored. If you run the command without the backup ID, you must manually update the configuration on the new device.

   During the RMA process, the following actions occur:
   - The device boots up in maintenance mode.
   - XCO updates the device ID for the connection details in the database.
   - Maintenance mode is initiated if the device is not already in maintenance mode.
   - XCO replays the backed-up configuration specified by the `config-backup-id` parameter of the **`efa inventory rma execute`** command.
   - XCO begins the drift reconcile process, which involves device discovery, device update, and fabric and tenant reconciliation. For more information, see Drift and Reconcile on page 60.

     Note
     If the RMA command fails during this stage, you can manually run the drift reconcile process from the CLI. If the RMA process fails for any other reason, restart the RMA process.

   - When drift reconcile is complete, the device is taken out of maintenance mode.
   - During the RMA process, XCO health checks are deactivated and RASlog does not trigger drift reconcile.
5. Install the HTTPS or OAuth2 certificate on the new device.

   Note
   The following conditions result in the RMA process failing:
   - If there is mismatch in device IP and credentials
   - If the device maintenance mode is not successful

   In both the conditions, reject the device replacement with the new device.

Replace a Faulty Device

You can use the XCO command line to replace a faulty device with a new device and maintain the configuration of the old device.

Before You Begin

- Ensure that periodic configuration backup is enabled on all devices that may need RMA. This prerequisite ensures that you have the latest configuration file to be used for recovery.
- Ensure that maintenance mode is enabled upon reboot on all devices.

You can replace the following SLX devices with the Extreme devices:

- SLX 9250 with Extreme 8720
- SLX 9150 with Extreme 8520
- SLX 9740 with Extreme 8820

About This Task

This procedure describes how to replace a faulty device as part of the Return Material Authorization (RMA) process. For more information, see Return Material Authorization on page 461.

Procedure

1. Obtain the configuration backup of the old device.

   For more information, see Configure Backup and Replay on page 460.

   ```
   # efa inventory config-backup execute --ip <ip-addr>
   ```

   This step generates a configuration backup ID that you use in step 5.
2. Replace the faulty device with the new device.
3. Ensure that the ports of the new device are administratively up and online.

   The ports on the new device must have the same connections as the old device. For example, if the old device Ethernet port1 and port2 went to port3 and port4 of another device, the new replacement device must have these exact same Ethernet port connections.
4. Configure the new device with the same management IP address and credentials as the old device.
5. Start the RMA process.

   ```
   # efa inventory rma execute --ip <ip-addr> --config-backup-id <id>
   ```

   Note

   As a best practice, run the command with the `--config-backup-id <id>` option so that the configuration is properly restored. If you run the command without the backup ID, you must manually update the configuration on the new device.

6. View the RMA history and detail.

   ```
   # efa inventory rma history --ip <ip-addr>
   # efa inventory rma detail –uuid <uuid>
   ```

7. View the drift reconcile history and detail.

   ```
   # efa inventory drift-reconcile history --device-ip <ip-addr>
   # efa inventory drift-reconcile detail --uuid <uuid>
   ```

8. Install the HTTPS or OAuth2 certificate on the new device.

   ```
   # efa certificates device install --ips <device-ip-addr> --certType [https|token]
   ```

9. (Optional) Delete the RMA record.

   ```
   # efa inventory rma delete --ip <ip-addr>
   ```

10. (Optional) Manually start the drift reconcile process if the RMA process fails during the drift reconcile stage.

    ```
    # efa inventory drift-reconcile execute --ip <ip-addr> --reconcile
    ```

SLX Device Configuration

Enable Maintenance Mode on SLX Devices

You can enable maintenance mode on the SLX devices that XCO manages.

About This Task

By default, XCO performs Drift and Reconcile actions on the SLX devices that enter into maintenance mode after reboot, taking those devices out of maintenance mode after successfully reconciling the configuration on them. For more information about Drift and Reconcile, see Drift and Reconcile on page 60.

You can enable maintenance mode on SLX devices without triggering Drift and Reconcile. Take the following steps.

Procedure

1. Disable syslog.

```
efa inventory device execute-cli --ip 10.18.120.187 --command "no logging syslog-
server 10.18.120.140 use-vrf mgmt-vrf" --config
```

2. Enable maintenance mode.

```
efa inventory device setting update --maint-mode-enable Yes --ip 10.18.120.187
```

   The device remains in maintenance mode until you disable the mode.

   If both `maint-mode-enable` and `maintenance-mode-enable-on-reboot` are set on the device, the Drift and Reconcile action is not triggered on device reboot.

3. Disable maintenance mode.
   a. Enable syslog.

```
efa inventory device execute-cli --ip 10.18.120.187 --command "logging syslog-
server 10.18.120.140 use-vrf mgmt-vrf" --config
```

   b. Run Drift and Reconcile.

```
efa inventory drift-reconcile execute --ip 10.18.120.187 -reconcile
```

   The Drift and Reconcile process takes the device out of maintenance mode.

Configure Physical Port Speed

You can configure the speed for receiving and transmitting data on a physical port.

About This Task

Follow this procedure to configure physical port speed.

Tip

In SLX-OS, you can use the **`show interface ethernet`** command to see the current speed setting for the Ethernet interfaces on your device.

You can change the port speed for one or more IP addresses or for a specified fabric. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Note

The **efa inventory device interface set-speed** command is an operational (or exec) command, not a configuration command. With operational commands, there is no configuration persistence, no drift identification, and no configuration reconciliation. You run operational commands as needed.

Procedure

Run the **efa inventory device interface set-speed** command.

This example sets the port speed on multiple IP addresses.

```
efa inventory device interface set-speed --ip 10.25.225.167,10.24.48.131,
10.24.51.135 --if-name 0/20-22  --speed 25gbps
+---------------+-----+------+----------------+------------+---------+--------+
|   DeviceIP    | ID  | Name | Interface Type | Port Speed | Result  | Reason |
+---------------+-----+------+----------------+------------+---------+--------+
| 10.25.225.167 | 9   | 0/21 | ethernet       | 25gbps     | Success |        |
+               +-----+------+----------------+------------+---------+--------+
|               | 89  | 0/20 | ethernet       | 25gbps     | Success |        |
+               +-----+------+----------------+------------+---------+--------+
|               | 1   | 0/22 | ethernet       | 25gbps     | Success |        |
+---------------+-----+------+----------------+------------+---------+--------+
| 10.24.51.135  | 16  | 0/21 | ethernet       | 25gbps     | Success |        |
+               +-----+------+----------------+------------+---------+--------+
|               | 86  | 0/20 | ethernet       | 25gbps     | Success |        |
+               +-----+------+----------------+------------+---------+--------+
|               | 48  | 0/22 | ethernet       | 25gbps     | Success |        |
+---------------+-----+------+----------------+------------+---------+--------+
| 10.24.48.131  | 142 | 0/20 | ethernet       | 25gbps     | Success |        |
+               +-----+------+----------------+------------+---------+--------+
|               | 110 | 0/21 | ethernet       | 25gbps     | Success |        |
+               +-----+------+----------------+------------+---------+--------+
|               | 148 | 0/22 | ethernet       | 25gbps     | Success |        |
+---------------+-----+------+----------------+------------+---------+--------+
Interface Details
--- Time Elapsed: 53.8631425s ---
```

This example sets the port speed for the specified fabric.

```
efa inventory device interface set-speed --fabric nc_no_vni
--if-name 0/20-22 --speed 25gbps
+--------------+-----+------+----------------+------------+---------+--------+
|   DeviceIP   | ID  | Name | Interface Type | Port Speed | Result  | Reason |
+--------------+-----+------+----------------+------------+---------+--------+
| 10.24.51.135 | 86  | 0/20 | ethernet       | 25gbps     | Success |        |
+              +-----+------+----------------+------------+---------+--------+
|              | 48  | 0/22 | ethernet       | 25gbps     | Success |        |
+              +-----+------+----------------+------------+---------+--------+
|              | 16  | 0/21 | ethernet       | 25gbps     | Success |        |
+--------------+-----+------+----------------+------------+---------+--------+
| 10.24.48.131 | 142 | 0/20 | ethernet       | 25gbps     | Success |        |
+              +-----+------+----------------+------------+---------+--------+
|              | 110 | 0/21 | ethernet       | 25gbps     | Success |        |
+              +-----+------+----------------+------------+---------+--------+
|              | 148 | 0/22 | ethernet       | 25gbps     | Success |        |
+--------------+-----+------+----------------+------------+---------+--------+
```

```
Interface Details
--- Time Elapsed: 36.9974805s ---
```

Configure Breakout Ports

You can break a port into multiple interfaces, such as breaking one 40G port into four 10G ports. You can also revert the breakout.

About This Task

In SLX-OS, you can use the **show running-config hardware** command to determine whether breakout mode is configured for a device.

You can break a port into the following modes: one 10g port, one 25g port, one 100g port, two 40g ports, two 50g ports, four 10g ports, and four 25g ports.

The breakout interfaces you create are identified by the name of the original interface followed by a suffix.

When you run revert a breakout, the breakout interfaces are deconfigured and deleted. The original Ethernet interface in the default configuration is created automatically.

You can configure breakout for one or more IP addresses or for a specified fabric. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Note
The **efa inventory device interface set-breakout** command is an operational (or exec) command, not a configuration command. With operational commands, there is no configuration persistence, no drift identification, and no configuration reconciliation. You run operational commands as needed.

Procedure

1. To break a port into multiple ports, run the **efa inventory device interface set-breakout** command.
   This example breaks three interfaces into four ports each.

```
efa inventory device interface set-breakout --ip 10.24.80.158
+--------------+----+-------+----------------+---------+--------+
|   DeviceIP   | ID | Name  | Interface Type | Result  | Reason |
+--------------+----+-------+----------------+---------+--------+
| 10.24.80.158 | 73 | 0/2:2 | ethernet       | Success |        |
+              +----+-------+----------------+---------+--------+
|              | 72 | 0/1:4 | ethernet       | Success |        |
+              +----+-------+----------------+---------+--------+
|              | 74 | 0/3:2 | ethernet       | Success |        |
+              +----+-------+----------------+---------+--------+
|              | 78 | 0/3:3 | ethernet       | Success |        |
+              +----+-------+----------------+---------+--------+
|              | 75 | 0/3:4 | ethernet       | Success |        |
+              +----+-------+----------------+---------+--------+
|              | 70 | 0/1:1 | ethernet       | Success |        |
+              +----+-------+----------------+---------+--------+
|              | 71 | 0/1:3 | ethernet       | Success |        |
```

```
+               +----+-------+----------------+---------+--------+
|               | 80 | 0/2:1 | ethernet       | Success |        |
+               +----+-------+----------------+---------+--------+
|               | 79 | 0/1:2 | ethernet       | Success |        |
+               +----+-------+----------------+---------+--------+
|               | 76 | 0/2:3 | ethernet       | Success |        |
+               +----+-------+----------------+---------+--------+
|               | 69 | 0/3:1 | ethernet       | Success |        |
+               +----+-------+----------------+---------+--------+
|               | 77 | 0/2:4 | ethernet       | Success |        |
+---------------+----+-------+----------------+---------+--------+
Interface Details
--- Time Elapsed: 48.3801684s ---
```

This example configures break out for the specified fabric.

```
efa inventory device interface set-breakout --if-name 0/19-20
--mode 4x25g --fabric fabric1
+---------------+-----+--------+----------------+---------+--------+
|   DeviceIP    | ID  |  Name  | Interface Type | Result  | Reason |
+---------------+-----+--------+----------------+---------+--------+
| 10.24.80.158 | 188 | 0/20:2 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 184 | 0/20:3 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 190 | 0/20:4 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 191 | 0/19:1 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 187 | 0/19:2 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 189 | 0/19:3 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 185 | 0/19:4 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 186 | 0/20:1 | ethernet       | Success |        |
+---------------+-----+--------+----------------+---------+--------+
| 10.24.80.159 | 196 | 0/19:2 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 194 | 0/19:3 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 197 | 0/19:4 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 199 | 0/20:1 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 193 | 0/20:2 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 195 | 0/20:3 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 192 | 0/20:4 | ethernet       | Success |        |
+               +-----+--------+----------------+---------+--------+
|               | 198 | 0/19:1 | ethernet       | Success |        |
+---------------+-----+--------+----------------+---------+--------+
Interface Details
--- Time Elapsed: 1m6.2210288s ---
```

2. To revert the breakout of multiple ports to the original configuration, run the **`efa inventory device interface unset-breakout`** command.

   This example removes breakout mode on multiple devices.

```
efa inventory device interface unset-breakout
--ip 10.24.80.158,10.24.80.159 --if-name 0/9-12
+---------------+--------------+----------------+----------------+---------+
|   DeviceIP    | Interface ID | Interface Name | Interface Type | Result  |
```

```
+--------------+--------------+----------------+----------------+---------+
| 10.24.80.158 | 248          | 0/10           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 250          | 0/11           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 249          | 0/12           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 247          | 0/9            | ethernet       | Success |
+--------------+--------------+----------------+----------------+---------+
| 10.24.80.159 | 252          | 0/10           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 254          | 0/11           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 253          | 0/12           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 251          | 0/9            | ethernet       | Success |
+--------------+--------------+----------------+----------------+---------+
Interface Details
--- Time Elapsed: 1m52.8562333s ---
```

This example removes breakout mode for the specified fabric name.

```
efa inventory device interface unset-breakout --if-name 0/19-20
--fabric fabric1
+--------------+--------------+----------------+----------------+---------+
|   DeviceIP   | Interface ID | Interface Name | Interface Type | Result  |
+--------------+--------------+----------------+----------------+---------+
| 10.24.80.159 | 279          | 0/19           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 280          | 0/20           | ethernet       | Success |
+--------------+--------------+----------------+----------------+---------+
| 10.24.80.158 | 281          | 0/19           | ethernet       | Success |
+              +--------------+----------------+----------------+---------+
|              | 282          | 0/20           | ethernet       | Success |
+--------------+--------------+----------------+----------------+---------+
Interface Details
--- Time Elapsed: 1m19.226463s ---
```

Configure MTU at Interface or System Level

You can configure the MTU (maximum transmission unit) at the system level or at the physical port level for Layer 2, IPv4, and IPv6.

About This Task

Follow this procedure to configure MTU at interface or system level.

Tip

In SLX-OS, you can use the **`show interface ethernet`** command to see the current MTU configuration for an interface.

You can change the MTU for one or more IP addresses or for a specified fabric. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Note

The **`efa inventory device interface set-mtu`** command is an operational (or exec) command, not a configuration command. With operational commands, there is no configuration persistence, no drift identification, and no configuration reconciliation. You run operational commands as needed.

Procedure

1. At the interface level, run the **`efa inventory device interface set-mtu`** command.

   This example configures the MTU on multiple IP addresses.

```
efa inventory device interface set-mtu --ip 10.25.225.167,10.24.48.131,
10.24.51.135 --if-name 0/20-22  --mtu 3600 --ip-mtu 3600 --ipv6-mtu 3600
+---------------+-----+------+----------------+------+--------+----------+---------
|   DeviceIP    | ID  | Name | Interface Type | MTU  | IP MTU | IPv6 MTU | Result
+---------------+-----+------+----------------+------+--------+----------+---------
| 10.25.225.167 | 9   | 0/21 | ethernet       | 3600 | 3600   | 3600     | Success
+               +-----+------+----------------+------+--------+----------+---------
|               | 89  | 0/20 | ethernet       | 3600 | 3600   | 3600     | Success
+               +-----+------+----------------+------+--------+----------+---------
|               | 1   | 0/22 | ethernet       | 3600 | 3600   | 3600     | Success
+---------------+-----+------+----------------+------+--------+----------+---------
| 10.24.48.131  | 142 | 0/20 | ethernet       | 3600 | 3600   | 3600     | Success
+               +-----+------+----------------+------+--------+----------+---------
|               | 148 | 0/22 | ethernet       | 3600 | 3600   | 3600     | Success
+               +-----+------+----------------+------+--------+----------+---------
|               | 110 | 0/21 | ethernet       | 3600 | 3600   | 3600     | Success
+---------------+-----+------+----------------+------+--------+----------+---------
| 10.24.51.135  | 16  | 0/21 | ethernet       | 3600 | 3600   | 3600     | Success
+               +-----+------+----------------+------+--------+----------+---------
|               | 48  | 0/22 | ethernet       | 3600 | 3600   | 3600     | Success
+               +-----+------+----------------+------+--------+----------+---------
|               | 86  | 0/20 | ethernet       | 3600 | 3600   | 3600     | Success
+---------------+-----+------+----------------+------+--------+----------+---------
Interface MTU Details
--- Time Elapsed: 59.5462548s ---
```

   This example configures the MTU for the specified fabric name.

```
efa inventory device interface set-mtu --fabric nc_no_vni --if-name 0/20-22,0/55-58
--mtu 3200 --ip-mtu 3200 --ipv6-mtu 3200
+-------+-----+--------+---------+-----+-----+-----+--------+--------------------------
|Device | ID  |  Name  |Interface| MTU | IP  |IPv6 | Result |       Reason
|IP     |     |        |Type     |     | MTU |MTU  |        |
+-------+-----+--------+---------+-----+-----+-----+--------+--------------------------
|10.24  | 16  | 0/21   |ethernet | 3200| 3200| 3200| Success|
|.51.135|     |        |         |     |     |     |        |
+       +-----+--------+---------+-----+-----+-----+--------+--------------------------
|       | 48  | 0/22   |ethernet | 3200| 3200| 3200| Success|
+       +-----+--------+---------+-----+-----+-----+--------+--------------------------
|       | 86  | 0/20   |ethernet | 3200| 3200| 3200| Success|
+       +-----+--------+---------+-----+-----+-----+--------+--------------------------
|       | 0   | 0/55-58|ethernet | 0   | 0   | 0   | Failed | Interfaces [0/55-58]
|       |     |        |         |     |     |     |        | do not exist for device
|       |     |        |         |     |     |     |        | IP 10.24.51.135
|       |     |        |         |     |     |     |        | Specify a valid interface
+-------+-----+--------+---------+-----+-----+-----+--------+--------------------------
```

```
|10.24  | 0   | 0/55-58|ethernet | 0   | 0   | 0   | Failed | Interfaces [0/55-58]
|.48.131|     |        |         |     |     |     |        | does not exist for
|       |     |        |         |     |     |     |        | device IP 10.24.51.131
|       |     |        |         |     |     |     |        | Specify a valid interface
+       +-----+--------+---------+-----+-----+-----+--------+-------------------------
|       | 110 | 0/21   |ethernet | 3200| 3200| 3200| Success|
+       +-----+--------+---------+-----+-----+-----+--------+-------------------------
|       | 148 | 0/22   |ethernet | 3200| 3200| 3200| Success|
+       +-----+--------+---------+-----+-----+-----+--------+-------------------------
|       | 142 | 0/20   |ethernet | 3200| 3200| 3200| Success|
+-------+-----+--------+---------+-----+-----+-----+--------+-------------------------
Interface MTU Details
--- Time Elapsed: 37.3021602s ---
```

2. At the interface level, run the **`efa inventory device interface unset-mtu`** command.

   This example unsets the MTU and IP-MTU from an interface.

```
efa inventory device interface unset-mtu  --ip 10.20.24.10 --if-name=0/6 --mtu -ip-mtu
+------------+----+------+----------+------+-----+------+--------+-------+
|   DeviceIP | ID | Name | Interface| MTU  | IP  | IPv6 | Result | Reason|
|            | ID |      | Type     | MTU  | MTU | MTU  |        |       |
+------------+----+------+----------+------+-----+------+--------+-------+
| 10.20.24.10| 451| 0/6  | ethernet | true | true| false| Success|       |
+------------+----+------+----------+------+-----+------+--------+-------+
Interface MTU Details
--- Time Elapsed: 6.469215243s ---
```

Note

Drift and Reconcile does not enforce the unset-mtu setting and enables an OOB (out-of-band) change to stay on the device.

Change the Admin Status of an Interface

You can bring an interface administratively up or down.

About This Task

Follow this procedure to change the admin status of an interface.

Tip

In SLX-OS, you can use the **`show interface ethernet`** command to see the status of the Ethernet interfaces on your device.

You can change the Admin Status for one or more IP addresses or for a specified fabric. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Note

The **`efa inventory device interface set-admin-state`** command is an operational (or exec) command, not a configuration command. With operational commands, there is no configuration persistence, no drift identification, and no configuration reconciliation. You run operational commands as needed.

Procedure

Run the **efa inventory device interface set-admin-state** command.

This example changes the Admin State on multiple IP addresses.

```
efa inventory device interface set-admin-state
--ip 10.25.225.167,10.24.48.131,10.24.51.135 --if-name 0/20-22 --state up

+---------------+-----+------+----------------+--------------+---------+--------+
|   DeviceIP    | ID  | Name | Interface Type | Admin Status | Result  | Reason |
+---------------+-----+------+----------------+--------------+---------+--------+
| 10.24.48.131  | 110 | 0/21 | ethernet       | up           | Success |        |
+               +-----+------+----------------+--------------+---------+--------+
|               | 142 | 0/20 | ethernet       | up           | Success |        |
+               +-----+------+----------------+--------------+---------+--------+
|               | 148 | 0/22 | ethernet       | up           | Success |        |
+---------------+-----+------+----------------+--------------+---------+--------+
| 10.24.51.135  | 16  | 0/21 | ethernet       | up           | Success |        |
+               +-----+------+----------------+--------------+---------+--------+
|               | 48  | 0/22 | ethernet       | up           | Success |        |
+               +-----+------+----------------+--------------+---------+--------+
|               | 86  | 0/20 | ethernet       | up           | Success |        |
+---------------+-----+------+----------------+--------------+---------+--------+
| 10.25.225.167 | 9   | 0/21 | ethernet       | up           | Success |        |
+               +-----+------+----------------+--------------+---------+--------+
|               | 89  | 0/20 | ethernet       | up           | Success |        |
+               +-----+------+----------------+--------------+---------+--------+
|               | 1   | 0/22 | ethernet       | up           | Success |        |
+---------------+-----+------+----------------+--------------+---------+--------+
Interface Details
--- Time Elapsed: 56.4964544s ---
```

This example changes the Admin Status for the specified fabric.

```
efa inventory device interface set-admin-state --fabric nc_no_vni
--if-name 0/20-22 --state up
+--------------+-----+------+----------------+--------------+---------+--------+
|   DeviceIP   | ID  | Name | Interface Type | Admin Status | Result  | Reason |
+--------------+-----+------+----------------+--------------+---------+--------+
| 10.24.51.135 | 48  | 0/22 | ethernet       | up           | Success |        |
+              +-----+------+----------------+--------------+---------+--------+
|              | 16  | 0/21 | ethernet       | up           | Success |        |
+              +-----+------+----------------+--------------+---------+--------+
|              | 86  | 0/20 | ethernet       | up           | Success |        |
+--------------+-----+------+----------------+--------------+---------+--------+
| 10.24.48.131 | 148 | 0/22 | ethernet       | up           | Success |        |
+              +-----+------+----------------+--------------+---------+--------+
|              | 142 | 0/20 | ethernet       | up           | Success |        |
+              +-----+------+----------------+--------------+---------+--------+
|              | 110 | 0/21 | ethernet       | up           | Success |        |
+--------------+-----+------+----------------+--------------+---------+--------+
Interface Details
--- Time Elapsed: 37.9236481s ---
```

Configure NTP at Device and Fabric Levels

Use the native commands to configure the NTP server on the SLX device. The configuration is persisted in the XCO database. DRC is supported.

About This Task

Follow this procedure to configure NTP server at device and fabric level.

Procedure

1. Run the **efa inventory device ntp server create** command to create an NTP server.

```
efa inventory device ntp server create ?
Flags:
      --ip string                Comma separated range of device IP addresses.
      --ntp-ip string            NTP server IP address
      --auth-key int             Authentication key ID. Values 1 to 65535
      --auth-key-name string     Key name
      --encryption-type string   Encryption type. Valid values are md5, sha1
      --trusted-key bool         Trusted key.
      --fabric string            fabric name
```

   Example:

```
efa inventory device ntp server create –ntp-ip 3.3.3.3 --auth-key 1 --auth-key-name
ntpsecret --encryption-type md5 –trusted-key  --ip 10.20.246.10

efa inventory device ntp server create –ntp-ip 3.3.3.3 --auth-key 1 --auth-key-name
ntpsecret --encryption-type md5 –trusted-key  --fabic clos_fabric
```

2. On the SLX device, verify the NTP configuration.

```
SLX# show running-config ntp
ntp authentication-key 1 md5 $9$750C7e0ayuI31YUga1Clmg== encryption-level 7
ntp authenticate
ntp server 3.3.3.3 key 1
```

3. Run the **efa inventory device ntp server delete** command to delete an NTP server.

```
efa inventory device ntp server delete ?
Flags:
      --ip string        Comma separated range of device IP addresses.
      --ntp-ip string    NTP server IP address
      --fabric string    fabric name
```

   Example:

```
efa inventory device ntp server delete --ntp-ip 3.3.3.3 --ip 10.20.246.10

efa inventory device ntp server delete --ntp-ip 3.3.3.3 --fabric clos_fabric
```

4. Run the **efa inventory device ntp server list** command to display the list of NTP servers configured using XCO.

```
efa inventory device ntp server list ?
Flags:
      --ip string        Comma separated range of device IP addresses.
      --fabric string    fabric name
```

   Example:

```
efa inventory device ntp server ntp server list --ip 10.20.246.10

efa inventory device ntp server ntp server list --fabric clos_fabric
```

5. Run the **efa inventory device ntp disable-server** command to disable SLX acting as an NTP servers to other clients. SLX cannot be an NTP server to other hosts.

```
efa inventory device ntp disable-server ?
Flags:
      --ip string        Comma separated range of device IP addresses.
      --enable           Disable ntp server. Valid values are yes/no.
```

```
      --list             List disable-server on devices.
      --fabric string    fabric name
```

Example

- Disable the NTP server on given device
```
efa inventory device ntp disable-server --enable yes --ip 10.20.246.10
```
- Enable the NTP server on given device
```
efa inventory device ntp disable-server --enable no --ip 10.20.246.10
```
- Disable the NTP server at fabric level
```
efa inventory device ntp disable-server --enable yes --fabric clos_fabric
```
- Enable the NTP server at fabric level
```
efa inventory device ntp disable-server --enable no --fabric clos_fabric
```
- List the NTP disable-server on given device
```
efa inventory device ntp disable-server --list --ip 10.20.246.10
```
- List the NTP disable-server on at fabric level
```
efa inventory device ntp disable-server --list --fabric clos_fabric
```

Configure RME on Interface

You can configure Redundant Management Ethernet (RME) on the SLX interface.

About This Task

Follow this procedure to set threshold monitor options.

Use the following CLIs to configure Redundant Management Ethernet (RME) on the SLX interface. This feature is supported only on 9150, 9250, 9740 SLX platforms. The configuration set by you is persisted in XCO database. DRC is supported:

- **efa inventory device interface set-fec**
- **efa inventory device interface set-link-error-disable**
- **efa inventory device interface set-snmp**
  - **efa inventory device snmp community create**
  - **efa inventory device snmp community delete**
  - **efa inventory device snmp community list**
  - **efa inventory device snmp user create**
  - **efa inventory device snmp user delete**
  - **efa inventory device snmp user list**
  - **efa inventory device snmp host create**
  - **efa inventory device snmp host delete**
  - **efa inventory device snmp host list**

Note

For more information about the commands, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*

Internally, XCO automatically sets the PPS (packets per second) value after RME is enabled. For SLX 9150 and 9250 devices, PPS is set to 8000.

On SLX 9740, the BPS (bits per second) is set to 20000 Kbps after RME is enabled.

Procedure

1. Run the **`efa inventory device interface redundant-management`** command.

```
efa inventory device interface redundant-management [flags]
Flags:
      --ip string        Comma separated range of device IP addresses.
      --if-name string   only one interface name. Example: 0/50
      --enable string    Valid values: true, false
            --- Time Elapsed: 9.610987ms ---
```

   Example:

```
efa inventory device interface redundant-management --ip 10.20.246.10 --if-name 0/17
--enable true
```

   - Run the following command to disable RME.

```
efa inventory device interface redundant-management --ip 10.20.246.10 --if-name
0/17 --enable false
```

2. Configure device on the SLX.

```
SLX# show running-config interface eth 0/17
interface Ethernet 0/17
 redundant-management enable
 no shutdown
!
```

Device Configuration Synchronization

During the first service boot after upgrade to XCO, XCO queries the SNMP and NTP configuration on the device. These configurations persist in the database, which is managed by XCO.

Breakout interfaces and interfaces that are in admin-state down, with non-auto speed or non-default MTU values, persist in the database and are managed by XCO

If you update the configuration, use the XCO CLI, not the SLX-OS CLI on the device. This ensures that the device configuration matches the XCO configuration.

SLX Configuration Backup

You can back up the device running the configuration that can be included as part of the existing XCO backup.

- Only one copy per device is included in a particular backup file.
- During a backup operation, no configuration changes are allowed through XCO.
- You can run a backup on demand or at an interval of your choosing.
- Any failure during the backup process is reported.

CLI Commands for Backups

Device backup is integrated with the existing system backup CLI.

Additional parameters allow you to specify either the fabric name or list of devices for which to run a backup.

Showing and Updating Backup Settings

There are commands for showing and updating backup settings.

Showing Backup Settings

`efa system settings show` lets you display all the system settings that have been configured.

Example

```
efa system settings show
+----------------------------+-------------+
|          SETTING           |   VALUE     |
+----------------------------+-------------+
| Max Backup File Limit      | 5           |
+----------------------------+-------------+
| Max Supportsave File Limit | 5           |
+----------------------------+-------------+
| Backup Schedule            | 0 0 * * *   |
+----------------------------+-------------+
| Remote Server IP           | 10.20.241.7 |
+----------------------------+-------------+
| Remote Server Username     | root        |
+----------------------------+-------------+
| Remote Server Password     | ****        |
+----------------------------+-------------+
| Remote Server Directory    | /root/vinod/|
+----------------------------+-------------+
| Remote Transfer Protocol      | scp      |
+-------------------------------+----------+
| Periodic Device Config Backup | yes      |
+-------------------------------+----------+
--- Time Elapsed: 831.836375ms ---
```

For more information, see `efa system settings show` in the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Updating Backup Settings

There are settings for configuring remote server details where the backup is copied. In case remote server details are missing, the backup is copied on the same server where XCO is installed, which is also the current behavior of the system backup.

Passwords are encrypted using an AES algorithm and stored in the database.

Remote server validation is performed to validate whether the provided details of a remote server are valid or reachable (if you enter only an IP address, the application checks in the database for the remaining parameters - if they are missing, then it is treated as an error). All four parameters (IP, username, password, and directory-path) are expected for validation, either from the user or the database.

Transfer of a backup archive on a remote server is done through the SCP protocol.

The `efa system settings update` command lets you make the updates.

Example

```
efa system settings update --remote-server-ip ip/ipv6 10.20.241.7
--remote-server-username root --remote-server-password pass --remote-server-directory /
root/vinod/
Setting Update Successful
--- Time Elapsed: 148.800033ms ---
```

Resetting System Backup Settings

You can reset the updated system backup settings to default values.

The `efa system settings reset` command lets you make the updates.

Example

```
efa system settings reset --max-backup-files
Reset System Settings is Successful
```

For more information, see `efa system settings reset` in the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Backup

To perform the backup, use the `efa system backup` command with options for specifying fabric or device details.

For more information, see `efa system backup` in the *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Backup Scenarios

The following backup scenarios showcase the use of different options with the `efa system backup` command.

Run system backup on remote

```
efa system backup --remote
Generating backup of EFA...
Backup Location: /var/log/efa/backup/EFA-3.1.0-110-2022-03-28T11-37-00.936.tar
--- Time Elapsed: 5.741750131s ---
```

Run system backup without device configuration backup

```
efa system backup --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
--- Time Elapsed: 5.741750131s ---
```

Run system backup by taking configuration backup of all devices that are part of the fabric specified

```
efa system backup --fabric default --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
--- Time Elapsed: 5.741750131s ---
```

Run system backup by taking configuration backup of all fabrics and its devices

```
efa system backup --fabric-all --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
--- Time Elapsed: 5.741750131s ---
```

Run system backup by taking configuration backup of all devices that are specified

```
efa system backup --device-ip 10.20.1.2,10.20.1.3,10.20.1.4 --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
--- Time Elapsed: 5.741750131s ---
```

Error message: Fabric does not exist

```
efa system backup --fabric default --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
Backup is partially success:
    Fabric does not exist
--- Time Elapsed: 5.741750131s ---
```

Error message: Device not found

```
efa system backup --device-ip 10.20.1.5,10.20.1.6 --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
Backup is partially success:
    Device 10.20.1.5 not found
--- Time Elapsed: 5.741750131s ---
```

Error message: Operation not allowed

```
efa system backup --device-ip 10.20.1.2,10.20.1.3 --remote
Generating backup of EFA...
Backup Location: root@10.20.241.7:/root/vinod/EFA--3.1.0-110-2022-03-28T11-37-00.936.tar
Backup is partially success:
    Devices [10.20.1.2] failed to get config backup as its locked for configuration
change by process [Firmware download].
--- Time Elapsed: 5.741750131s ---
```

Backup Schedule

Existing features of the wider backup schedule work here, too. Additionally, the backup gets SLX configuration backup of all those devices that are associated with a valid fabric.

XCO Native Support for SLX Threshold Monitor Settings

Set Threshold Monitor Options

The threshold settings trigger events from SLX when the threshold limits are met. This feature is supported on SLX 20.4.2 and above.

About This Task

Follow this procedure to set threshold monitor options.

Table 21: Drift Reconcile & Idempotency Support

Identify Drift	Reconcile configuration	Idempotency
Yes	Yes	Yes

Procedure

Run the following command:

```
efa inventory device threshold-monitor set [ --ip device-ips | --fabric fabric-name |
--type { cpu | memory | bfd-session | lif | mac-table | vxlan-tunnel } | --actions
{ all none raslog snmp loginfo } | --high-limit percentage | --low-limit percentage |
--interval seconds | --retry retry-interval | --count max-events |
```

Parameters

```
--ip device-ips
Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.
--fabric fabric-name
Specify the name of the fabric.
--type type
Specify the type of threshold being configured. Valid types are cpu, memory, bfd-session,
lif, mac-table, and vxlan-tunnel
--actions actions
Specifies the action to take upon exceeding the high-limit or low-limit. Valid actions
are all, none, raslog, snmp, and loginfo. Loginfo is only valid for type cpu/memory.
[all] Both RASLOG and SNMP trap will be generated. [none] No action will be taken.
[raslog] RASLOG will be generated. [snmp] SNMP trap will be generated. [loginfo]
Diagnostic data collection along with RASLOG.
--high-limit percentage
Specifies the high limit threshold percentage. Valid values for type cpu is 0-95 with
Default 50. Valid values for type memory is 0-95 with Default 70.  Valid values for types
bfd-session/lif/mac-table/vxlan-tunnel is 10-100 with Default 90.
--low-limit percentage
Specifies the low limit threshold percentage. Not applicable for type cpu/memory.  Valid
values for types bfd-session/lif/mac-table/vxlan-tunnel is 10-99 with Default 70.
--count max-events
Specifies the maximum number of event actions that can be taken in an interval. See
interval. Valid values are 1-60 with Default 4. Valid for types bfd-session/lif/mac-table/
vxlan-tunnel.
--interval seconds
Specifies the interval for the threshold type.  For type cpu/memory it specifies the
polling interval, in seconds after which a sample will be taken. Valid values are 10-3600
seconds with default of 120 seconds.  For types bfd-session/lif/mac-table/vxlan-tunnel it
specifies the time interval, in seconds where a maximum number (count) of actions that
can be taken. Valid values are 60-900 with Default 120.
--retry retry-interval
Specifies the number of polling interval retries that will be attempted (see interval)
```

```
before desired action is taken.  Only applicable for –type cpu/memory. Valid values are
1-100. The default retries is 3.
```

Example

This example shows the current inventory threshold-monitor settings for the specified device.

```
efa inventory device threshold-monitor list --ip 10.10.10.75
+--------+-------------+--------+-----+------+------+---------+------+-----------+
|IP      |    Type     |Actions |High |Low   |Count |Interval |Retry |AppState   |
|Address |             |        |Limit|Limit |      |         |      |           |
+--------+-------------+--------+-----+------+------+---------+------+-----------+
|10.10   |bfd-session  |snmp    |80   | 50   |3     |60       |      |cfg-in-sync|
+.10.75  +-------------+--------+-----+------+------+---------+------+-----------+
|        |vxlan-tunnel |        |80   |      |6     |         |      |cfg-in-sync|
+        +-------------+--------+-----+------+------+---------+------+-----------+
|        |lif          |snmp    |80   | 50   |3     |60       |      |cfg-in-sync|
+        +-------------+--------+-----+------+------+---------+------+-----------+
|        |mac-table    |snmp    |80   | 50   |3     |60       |      |cfg-in-sync|
+        +-------------+--------+-----+------+------+---------+------+-----------+
|        |cpu          |all     |80   |      |      |60       |3     |cfg-in-sync|
+        +-------------+--------+-----+------+------+---------+------+-----------+
|        |memory       |raslog  |80   |      |      |60       |3     |cfg-in-sync|
+--------+-------------+--------+-----+------+------+---------+------+-----------+
```

Example

This example sets a CPU threshold-monitor to trigger an SNMP event when the CPU threshold exceeds 80%. To trigger an event, ensure that the threshold exceeds 3 times (retry) with a polling interval of 60 seconds. It takes a total of 3 minutes before the notification is sent.

```
efa inventory device threshold-monitor set –type cpu --ip 10.10.10.153-154 –actions
snmp --high-limit 80 interval 60 retry 3
```

This example sets a mac-table threshold-monitor to trigger an SNMP event when the high-limit threshold exceeds 80% or the low-limit threshold exceeds 50%. You can trigger a maximum of 3 (count) SNMP events over the 60 second interval for the threshold exceeding 80% (high) or 50% (low) limits.

```
efa inventory device threshold-monitor set –type mac-table --fabric fabric1 –actions
snmp --high-limit 80 –low-limit 50 interval 60 count 3
```

Unset Threshold Monitor Settings

You can unset threshold monitor settings.

About This Task

Use this procedure to unset threshold monitor settings.

Procedure

Run the following command:

```
efa inventory device threshold-monitor unset [ --ip device-ip --fabric fabric-name --type
{ cpu | memory | bfd-session | lif | mac-table | vxlan-tunnel }]
```

Parameters

```
--ip device-ips
Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.
--fabric fabric-name
Specify the name of the fabric.
--type type
Specify the type of threshold being unconfigured. Valid types are cpu, memory, bfd-
session, lif, mac-table, and vxlan-tunnel
```

Example

This example unsets the current inventory threshold-monitor settings for the specified device.

```
efa inventory device threshold-monitor unset --ip 10.10.10.75-76
```

This example unsets a single inventory threshold-monitor setting for the specified device.

```
efa inventory device threshold-monitor unset --ip 10.10.10.75 --type cpu
```

Display Threshold Monitor Settings

You can use the CLI commands to display threshold monitor settings which are part of the device settings..

About This Task

Use this procedure to display threshold monitor settings.

Procedure

Run the following command:

```
efa inventory device threshold-monitor list --ip device-ips --fabric fabric-name
```

Parameters

```
--ip device-ips
Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.
--fabric fabric-name
Specify the name of the fabric.
```

Example

This example shows the current inventory threshold-monitor settings for the specified device.

```
efa inventory device threshold-monitor list --ip 10.10.10.75
+--------+-------------+--------+------+------+------+---------+------+-----------+
|IP      |     Type    | Actions| High | Low  | Count| Interval| Retry| AppState  |
|Address |             |        | Limit| Limit|      |         |      |           |
+--------+-------------+--------+------+------+------+---------+------+-----------+
|10.10   | bfd-session | snmp   | 80   | 50   | 3    | 60      |      |cfg-in-sync|
|.10.75  |             |        |      |      |      |         |      |           |
+        +-------------+--------+------+------+------+---------+------+-----------+
|        | vxlan-tunnel|        | 80   |      | 6    |         |      |cfg-in-sync|
+        +-------------+--------+------+------+------+---------+------+-----------+
|        | lif         | snmp   | 80   | 50   | 3    | 60      |      |cfg-in-sync|
+        +-------------+--------+------+------+------+---------+------+-----------+
```

```
|        | mac-table   | snmp   | 80   | 50   | 3    | 60      |      |cfg-in-sync|
+        +-------------+--------+------+------+------+---------+------+-----------+
|        | cpu         | all    | 80   |      |      | 60      | 3    |cfg-in-sync|
+        +-------------+--------+------+------+------+---------+------+-----------+
|        | memory      | raslog | 80   |      |      | 60      | 3    |cfg-in-sync|
+--------+-------------+--------+------+------+------+---------+------+-----------+
```

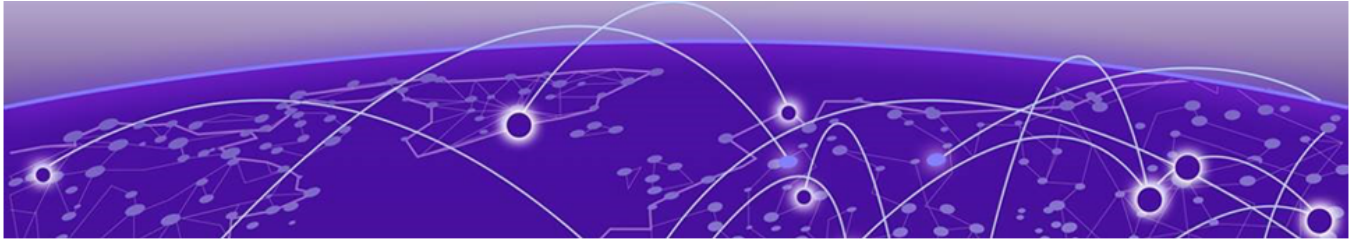

XCO Event Management

RASlog Service

The RASlog Service is aware of all devices that are registered with the services in XCO and processes events only from those devices. Messages from other devices are dropped.

The RASlog Service performs the following functions:

- Acts as a syslog server to process syslog messages from devices
- Acts as an SNMP trap receiver to process traps from devices

With the RASlog Service, XCO receives events from network devices and the Inventory service learns of relevant changes. The Inventory Service can fetch the current state of network topology and update Fabric and Tenant services.

RASlog Operations

XCO is registered as a syslog recipient on the devices as part of the device registration. If there are any changes to the link after fabric or tenant formation, the RASlog service receives the syslog message.

The sequence of RASlog operations is as follows:

1. The RASlog Service processes the syslog message and notifies all services through message-bus.
2. The Inventory Service receives the RASlog Service message and updates relevant asset details in the database.
3. The Inventory Service notifies Fabric and Tenant Services of any changes in the configurations.
4. Fabric and Tenant Services review the state changes and display information about any pending configurations.

   You can choose to update fabric or tenants for the current state.

5. When a device is deleted from the Inventory Service, XCO is unregistered as a syslog recipient from the device. If unregistration of XCO fails, deletion still proceeds.
6. The RASlog Service listens to Device Registration and Device Deletion messages to ensure that messages from registered devices are not dropped.

Notification Service

Overview

Notification service notifies the external entities about the events and alerts that occur on XCO and XCO managed devices.

Device events are derived from the syslog events that are received from the devices that are managed by XCO.

Alerts are notifications that XCO services send for unexpected conditions, such as the following:

- Loss of switch connectivity
- Failure to configure the fabric, tenant, or endpoint group (EPG) on the device
- Failure to perform operations such as port up or port down, set speeds, and breakout mode
- Firmware download failure
- Devices exiting maintenance mode
- Certificate expiry, expired alerts
- Storage threshold alerts

Task notifications are based on user-driven or timer-based operations, such as the following:

- Registering or updating a device
- Device timer collection completed
- Adding devices to a fabric
- Creating, updating, or deleting a fabric
- Creating, updating, or deleting a tenant
- Creating, updating, or deleting an endpoint group

Alarm notifications are sent out when alarms are raised or cleared, and when severities are updated.

Notification Methods

XCO supports two methods of notification: HTTPS webhook and syslog (using Reliable Event Logging Protocol [RELP] over Transport Layer Security [TLS]). The format of the notifications is the same for both methods. You can configure one or both methods.

Webhook

This REST API-based method is a POST operation. The notification payload is in the body of the HTTPS call. Use the **`efa notification subscribers add-https`** command to register a subscriber for this method of notification.

Syslog over RELP

In this client-server method, the client initiates the connection and the server listens. In this scenario, the client is the Notification service and the server is the remote system where syslog is configured to work with RELP. Any external server that is configured with RELP can be registered as a subscriber to XCO notifications.

When RELP is configured with TLS, XCO must be installed in secure mode. For more information, see the "XCO Installation Modes" topic in the *ExtremeCloud Orchestrator Deployment Guide, 3.2.1*.

Communication from SLX devices occurs over TLS. The certificates required for SLX devices to work with secure syslog are generated when the devices are registered.

Use the **`efa notification subscribers add-syslog-relp`** command to register a subscriber for this method of notification.

Note

Notification service will be disabled and will not send any alerts or events to those subscribers that fail to send messages due to connectivity issues. There will be a five minute periodic timer to verify the connectivity status of the failed subscribers.

EFA-008000 is a special event which will be received by the subscribers during the verification of connectivity between XCO notification service and subscribers. Subscribers receiving this message are indicated that there was a prior connectivity issue and it is now resolved.

Notification Types

XCO notifies the subscribers on the following event or alert types:

- DEVICE_EVENTS - Syslog or RASlog events generated by the devices
- APP_EVENTS - Task or user events generated by XCO
- APP_ALERTS – Alerts are generated by XCO. Prior to EFA 3.1.0, alerts were in non-standard or legacy format. EFA 3.1.0 introduces fault-alerts which has more alert specific data.
- APP_ALARMS – Alarms are generated by XCO.

By default, the subscribers receive all the event or alert types and users have the option to configure the filters during registration of subscribers. The filter option is applicable to both HTTPS and RELP subscribers.

Note

- Existing subscribers (registered in EFA 3.0.0 or earlier) receive all notification alert or event types when upgraded to EFA 3.1.0.
- For modifying the filter values, users must unsubscribe or delete and re-add it. EFA 3.1.0 does not support updating filter value.
- Webhook subscribers receive only fault-alerts introduced in EFA 3.1.0.
- Syslog subscribers receive either legacy alerts or fault-alerts based on the RFC5424 flag.

Notification Sub-Filtering

XCO provides an additional filtering capability of logging stream, including sub-filtering of SLX messages by message type and filtering of XCO and SLX messages by minimum severity level.

Sub-filtering of DEVICE_EVENTS notifications are only supported on SLX devices.

Terminology

The following table describes the messaging terminology used for notifying events in XCO system.

Abbreviation	Expansion
SLX Messages	SLX log messages, specifically, RASLog and AuditLog messages.
RASLog Messages	RASLog - Reliability, Availability and Serviceability (RAS) log messages.
AuditLog Messages	Event auditing messages to support post-event audits and problem determination. It is based on high-frequency events of certain types, such as security violations, firmware downloads, and configuration. Event auditing is broken down into the following three types: 1. Configuration 2. Firmware 3. Security
Configuration	Audit all configuration changes. A subtype of AUDIT.
Firmware	Audit the events that occur during the firmware download process. A subtype of AUDIT.

Abbreviation	Expansion
Security	Audit any user-initiated security event for all management interfaces. A subtype of AUDIT.
Severity	Message severity type and order (from most important to least): 1. CRITICAL 2. ERROR 3. WARNING 4. INFO

Additional Notification Filtering

A sub-filtering capability of the logging streams includes filtering of device, XCO, and SLX events or alert types by minimum severity level.

Device Event Sub-Filtering

Device Events send two types of messages: RASLOG and AUDIT. To refine your filter, use a combination of the keywords, such as raslog, audit-configuration, audit-firmware, and audit-security.

The sub-filtering is only applicable for SLX devices.

Event and Alert Sub-Filtering by Minimum Severity Level

You can reduce the filtering and device event sub-filtering notifications by filtering XCO and SLX events or alert types by minimum severity level. If you provide Info or no sub-filter value for the sub-filter, then no filtering will be done. A higher level severity value, such as Critical, Error or Warning results in filtering out all the messages of lower severity.

The device alerts and alarms messages of major or critical severity are not filtered because they are at or above the highest minimum-severity level. The device alerts and alarms messages of minor severity are treated as severity level of error for sub-filtering.

The following table describes the use of commands for filtering device events:

Commands	Description
`–device-event “audit-configuration”,“audit-firmware", “audit-security"`	Receives all the audit messages but filter out all “raslog” messages for applicable device events.
`–device-event “audit-configuration”,“audit-security”`	Receives no raslog notifications and only configuration and security-related audit notifications.
`--minimum-severity-subfilter “warning”`	Filters out all “Info” messages for applicable device events, app events, alerts, and alarms.

Sub-Filter CLI

You can enable subscribers with all the notification types of device events and alerts.

Use the following command to a register a new subscriber to the Notification service with an HTTPS webhook:

```
(efa:root)XCO-Server# efa notification subscribers add-https
Register a new https webhook subscriber to the notification service.
```

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 3.2.1*

1. The following example enables webhook subscriber only with device event audit configuration:

```
#efa notification subscribers add-https --url https://127.0.0.1:5000 -username
jarvis --password vision --insecure --filter DEVICE_EVENTS --device-event audit-
firmware,audit-security,audit-configuration
Successfully registered subscriber.

+-----------+-----------------------------------------------------+
| attribute | value                                               |
+-----------+-----------------------------------------------------+
| id        | 18                                                  |
+-----------+-----------------------------------------------------+
| handler   | http                                                |
+-----------+-----------------------------------------------------+
| endpoint  | https://127.0.0.1:5000                              |
+-----------+-----------------------------------------------------+
| config    | {"cacert":"","filters":["DEVICE_EVENTS"],           |
|           |”device-event”:[”audit-firmware”,audit-security",    |
|           | ”audit-configuration”],"insecure":true,"password":|
|           |"vision","username":"jarvis"}                        |
+-----------+-----------------------------------------------------+
Notification Subscriber ID=18
--- Time Elapsed: 2.148580069s ---
```

2. The following example enables warning or higher notification of app alerts and device events, and only device events with audit-security and audit-configuration:

```
#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure
--filter APP_ALERTS,DEVICE_EVENTS --device-event audit-security,audit-configuration --
minimum-severity warning
Successfully registered subscriber.

+-----------+-----------------------------------------------------------+
| attribute | value                                                     |
+-----------+-----------------------------------------------------------+
| id        | 19                                                        |
+-----------+-----------------------------------------------------------+
| handler   | relp                                                      |
+-----------+-----------------------------------------------------------+
| endpoint  | 127.0.0.1:1601                                            |
+-----------+-----------------------------------------------------------+
| config    | {"cacert":"","conn-timeout":10,"filters":["APP_ALERTS",|
|           | "DEVICE_EVENTS"],”device-event”:[”audit-security”,      |
|           | } ”audit-configuration”],”minimum-severity”,”warning”,    |
|           | "insecure":true}                                          |
+-----------+-----------------------------------------------------------+
Notification Subscriber ID=19
--- Time Elapsed: 2.172557260s ---
```

3. The following example enables all notification types on syslog subscriber of severity error or higher:

```
#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure --
minimum-severity error
Successfully registered subscriber.

+-----------+---------------------------------------------+
| attribute | value                                       |
+-----------+---------------------------------------------+
| id        | 20                                          |
+-----------+---------------------------------------------+
| handler   | relp                                        |
+-----------+---------------------------------------------+
| endpoint  | 127.0.0.1:1601                              |
+-----------+---------------------------------------------+
| config    | {"cacert":"","conn-timeout":10,"filters":[],|
|           | ”minimum-severity”,”error”,"insecure":true} |
+-----------+---------------------------------------------+
Notification Subscriber ID=20
--- Time Elapsed: 2.042797881s ---
```

Sub-Filter Options during XCO Upgrade

XCO upgrade (for example, from EFA 3.1.0 to XCO 3.2.0) with log streaming contains default sub-filters which is equivalent to no sub-filter.

- For a device event, the default is equivalent to a sub-filter with raslog, audit-configuration, audit-firmware, and audit-security.
- For minimum severity, the default is equivalent to a sub-filter with info.

When you upgrade from EFA 3.1.0 to XCO 3.2.0, changes to the external webhook or syslog server are not necessary.

Webhooks Payload

Webhooks payload is a key-value pair which can hold information for all notification types. Some fields are common and some are applicable only to particular notification types. Webhooks send only the fault alerts introduced in EFA 3.1.0.

The following table summarizes the supported key-value pairs:

Payload Fields	APP_EVENTS (Task Events)	DEVICE_EVENTS (SLX Raslog Events)	APP_ALERTS (Fault-Alerts Only)	APP_ALARMS (Fault Alarms Only)
type	✓ (Task)	✓ (Event)	✓ (Alert)	✓ (Alarm)
timestamp	✓	✓	✓	✓
severity	✓	✓	✓	✓
message	✓	✓	✓	✓
application	✓	✓	✓	✓
source_ip	×	✓	✓	✓

Payload Fields	APP_EVENTS (Task Events)	DEVICE_EVENTS (SLX Raslog Events)	APP_ALERTS (Fault-Alerts Only)	APP_ALARMS (Fault Alarms Only)
device_ip	✓	✓	×	×
username	✓	✓	×	×
message_id	×	✓	×	×
hostname	×	✓	✓	✓
logtype	×	✓	×	×
task	✓	×	×	×
scope	✓	×	×	×
status	✓	×	×	×
sequence_id	×	×	✓	✓
alert_id	×	×	✓	×
alarm_id	×	×	×	✓
resource	×	×	✓	✓
alarm_type	×	×	✓	✓
alarm_cause	×	×	✓	✓
alert_data	×	×	✓ (This will have nested key-value pairs with alert specific data)	×
alarm_data	×	×	×	✓ (This will have nested key-value pairs with alarm specific data)

Syslog Subscribers Message Format

EFA 3.1.0 introduces RFC5424 format for all notification types. Using RFC5424 flag at the subscriber level, syslog subscribers can choose the message format.

The following table summarizes the effect of the RFC5424 flag on syslog subscribers:

RFC-5424	Device Events	App Events	App Alerts
Enabled	RFC-5424 format	RFC-5424 format	RFC-5424 format (Only fault-alerts will be sent)
Disabled	Legacy format	Legacy format	Legacy format (Only legacy alerts which exists prior to 3.1.0 will be sent)

Note

1. For existing subscribers (registered in EFA 3.0.0 or earlier) RFC5424 flag is disabled after upgrading to EFA 3.1.0.
2. For modifying the RFC5424 flag, unsubscribe or delete and re-add it. EFA 3.1.0 does not support updating the RFC5424 flag.

CLI changes

```
(efa:root)root@pasu-dev-server:~/build/efa# efa notification subscribers    add-syslog-
relp --help
Register a new RELP syslog subscriber to the notification service.

Usage:
  efa notification subscribers add-syslog-relp [flags]

Flags:
      --address string     Address for syslog server in the format host:port (Required),
Default port: 514
      --insecure           Perform insecure SSL connection and transfers. (Optional)
      --cacert string      Local path to the cacert pem file for SSL verification.
(Optional, required if not insecure)
      --conn-timeout int   Timeout to open a connection to the server (Optional) (default
10)
      --filter strings     Comma separated filter values. Possible values are
"DEVICE_EVENTS" - RAS/syslog events from devices, "APP_ALERTS" - fault alerts from
application, "APP_EVENTS" - task events from application. If no filters are provided it
means all types. E.g. --filters DEVICE_EVENTS,APP_ALERTS,APP_EVENTS. (Optional)
      --rfc5424            Enable RFC5424 message format for syslog subscribers
(Optional) (default: non-RFC5424 format)
  --- Time Elapsed: 2.756476ms ---
```

The following example enables RFC-5424 format:

```
#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601  --insecure --
rfc5424
Successfully registered subscriber.

+----------+-------------------------------------------------------------------------------+
|attribute |value                                                                          |
+----------+-------------------------------------------------------------------------------+
|id        |7                                                                              |
+----------+-------------------------------------------------------------------------------+
|handler   |relp                                                                           |
+----------+-------------------------------------------------------------------------------+
|endpoint  |134.141.21.190:1601                                                            |
+----------+-------------------------------------------------------------------------------+
|config    |{"cacert":"","conn-timeout":10,"filters":[],"insecure":true,"rfc5424":true}|
```

```
+----------+-------------------------------------------------------------------------------+
Notification Subscriber ID=7
```

App Events RFC-5424 Format

This provides the common fields of the APP_EVENTS object that would be sent over the Syslog channel.

Field	SD-ID (Structured Data ID)	Example	Description
<###>	*N/A*	*190* =(**23** * 8) + 6	Priority Value: (*Syslog Classifier* * 8) + *Syslog Informational message* *Syslog Classifier*: 23 Local7 *Syslog Severity*: 6 Informational: informational messages
Version	*N/A*	*1*	Version of syslog message
Timestamp	*N/A*	*2003-10-11T22:14:15.003Z*	Timestamp of syslog message
Hostname	*N/A*	*xco.machine.com*	Hostname of XCO
App Name	*N/A*	*XCO-fabric*	Application generating syslog alerts. Possible values • *XCO-inventory* • *XCO-evm* • *XCO-policy* • XCO-ts • XCO-fabric
Proc ID	*N/A*	-	*Process ID*
Msg ID	*N/A*	-	
IP	*origin*	*10.20.30.40*	*IP address (of XCO host)*
Enterprise ID	*origin*	*1916*	*Extreme Networks Enterprise ID*
Software	*origin*	*XCO*	*Software Name (of XCO host)*
SW Version	*origin*	*3.1.0*	*Software Version (of EFA host)*
Taskname	*log@1916*	*XCO-000001*	*Task name ranges are defined as follows:* *Fabric – XCO-000001 to XCO-001000* *Tenant – XCO-001001 to XCO-002000* *Inventory – XCO-002001 to XCO-003000* *Policy – XCO-003001 to XCO-003059*

Field	SD-ID (Structured Data ID)	Example	Description
Scope	*log@1916*	*user*	*Scope of the task "user" or "system". Currently only user level scope is supported.*
Status	*log@1916*	*succeeded*	*Status of the task "started", "succeeded" or "failed"*
DeviceIP	*log@1916*	*""*	*Device IP involved in the user task*
Username	*log@1916*	*admin*	*User name*
Severity	*log@1916*	*Info*	*Severity is always "info"*
BOMText	*N/A*		*(Byte Order Mask) Textual description of the Alert*

Map APP_EVENTS to RELP/Syslog fields (RFC-5424)

```
<190>1 2022-10-10T21:29:45-07:00 pasu-dev-server EFA-ts - -
[origin ip="10.20.241.27" enterpriseId="1916" software="EFA" swVersion="3.1.0 "]
[log@1916 taskname="EFA-001002" scope="user" status="succeeded" deviceip=""
username="root" severity="info"]
BOM Tenant create request success :request={"name":"ts"}
```

Device Events RFC-5424 Format

This provides the common fields of the DEVICE_EVENTS object that would be sent over the Syslog channel.

Field	SD-ID (Structured Data ID)	Example	Description
<###>	*N/A*	*190* =(**23** * 8) + 6	Priority Value: (*Syslog Classifier* * 8) + *Syslog Informational message Syslog Classifier*: 23 Local7 *Syslog Severity*: 6 Informational: informational messages
Version	*N/A*	*1*	Version of syslog message
Timestamp	*N/A*	*2022-10-12T09:54:28.506827-07:00*	Timestamp of syslog message
Hostname	*N/A*	*GE-SH*	Hostname of XCO
App Name	*N/A*	-	Application generating syslog alerts. Possible values
Proc ID	*N/A*	-	*Process ID*
Msg ID	*N/A*	*DCM-1116*	*Raslog ID of the device event*

Field	SD-ID (Structured Data ID)	Example	Description
Sequence ID	*meta*	*49051*	*Tracks the sequence in which messages are submitted to the syslog transport per device.*
IP	*origin*	*10.20.30.40*	*IP address (of the Device)*
Enterprise ID	*origin*	*1588*	*Device's Enterprise ID*
BOMText	*N/A*	*System initialization is complete. SLX-OS is ready to handle all commands.*	*(Byte Order Mask) Textual description of the Alert*

Map DEVICE_EVENTS to RELP/Syslog fields (RFC-5424)

```
<190>1 2022-10-12T09:54:28.506827-07:00 GE-SH - - DCM-1116
[meta sequenceId="49051"]
[origin ip="10.24.0.56" enterpriseId="1588"]
[log@1588 value="RASLOG"]
[timestamp@1588 value="2022-10-12T16:54:24.395333"]
[msgid@1588 value="DCM-1116"]
[attr@1588 value=" WWN 10:00:00:04:96:b8:37:5b"]
[severity@1588 value="INFO"]
[swname@1588 value="SLX9740-80C"]
BOMSystem initialization is complete. SLX-OS is ready to handle all commands.
```

XCO as SNMP Proxy

Simple Network Management Protocol (SNMP) traps are alert messages sent from a remote SNMP-enabled device to a central collector, the SNMP Manager. Trap messages are the main form of communication between SNMP monitoring tools – an SNMP Agent and an SNMP Manager.

XCO acts as the SNMP Manager for all the SLX devices and agents and receives the traps from all the devices in its inventory. Once you register an SLX device with XCO, XCO automatically configures the SLX device to send v3 traps to XCO.

XCO acts as an SNMP proxy for all the SNMP v2 and v3 traps received from the SLX devices, forwarding them onto an external trap receiver, if there is one.

- XCO subscribes to be a v3 trap receiver with a predefined v3 user name, authentication key, and privacy key.
- If you set up XCO to be a v2c trap receiver, you must provide a community string.

During an update operation, XCO verifies that it is still registered to receive traps from the SLX devices. If a device is unregistered from XCO, the SNMP configuration on the device is updated to no longer send traps to the XCO IP address.

Commands for configuring SNMP

The following commands are available for configuring SNMP on the SLX device. The configuration you set is persisted in the XCO database. DRC is supported.

- `efa inventory device snmp community create`
- `efa inventory device snmp community delete`
- `efa inventory device snmp community list`
- `efa inventory device snmp user create`
- `efa inventory device snmp user delete`
- `efa inventory device snmp user list`
- `efa inventory device snmp host create`
- `efa inventory device snmp host delete`
- `efa inventory device snmp host list`
- `efa inventory device snmp view create`
- `efa inventory device snmp view delete`
- `efa inventory device snmp view list`

For more information about these commands, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Notes

- The device IP address is the one included in `SNMP-COMMUNITY-MIB::snmpTrapAddress.0`. It is not the XCO IP address.
- XCO forwards all received traps. In other words, no trap is filtered out.
- Port 162 on the host where XCO is installed must be available. During a fresh installation, the port availability is checked and the installer returns an error if the port is not available. However, during an upgrade from a previous version of XCO, you must ensure that the port is free.

For more information about SLX-OS MIBs, see the *Extreme SLX-OS MIB Reference* for your version of SLX-OS.

Limitations

- A maximum of four trap subscribers is supported.
- V2c and v3 SNMP subscribers are not validated.
- Only traps generated by SLX devices are forwarded. Alerts and alarms from XCO itself are not forwarded.
- Only traps are forwarded. Current XCO tasks or alerts and syslog messages are not forwarded as traps.
- SNMP Informs are not supported.
- There is no in-band support for trap forwarding.
- The Drift and Reconcile process does not show a drift in device configuration for SNMP v3 trap configuration that XCO has pushed. However, every time the device

update operation runs, XCO checks if the device is configured to send traps to XCO and if not, pushes the configuration again.

- For a multi-node deployment during failover of the active node, some traps might be missed while the SNMP service is bootstrapping on the new active node. There is no loss of traps if the standby node goes down.

Migration of existing switch configuration

When you boot the service for the first time after upgrade, any SNMP and NTP configuration on the switch are queried and persisted in the database and managed by XCO.

Similarly, any breakout interfaces or interfaces that have status admin-state DOWN and have a non-auto speed or non-default MTU value are persisted in the database and managed by XCO.

If you have additional updates to make to these configurations, you must make them manually using the XCO commands only.

If these configurations are updated using the SLX commands directly on the switch (meaning, not by using the XCO CLI), they are considered as drifted and are reconciled.

gosnmp-service

The gosnmp-service is responsible for persisting the trap subscribers, receiving the SNMP traps, and forwarding them to the subscribers.

The service is stateless, so no historical data (that is, previously received traps) is persisted.

For high availability deployment, the service runs in active-active mode, however, since the VIP is bound to one host at a time, the pod running on the active node receives the traps. On failover, the standby node takes over and the SNMP service running on that node forwards the traps.

You may have multiple IP subnets configured to access XCO. In such a case, XCO creates multiple subinterfaces under the management interface to which XCO is bound. XCO does not determine which interface sends out the trap, syslog or webhook. The administrator is responsible for configuring a route to the recipient. If one is found, the server sends out the trap. For more information, see Multiple Management IP Networks on page 76.

Configure SNMP View and Destination UDP Port

SNMP view is a group of MIB OIDs that limits viewing and configuring access within SNMP. SNMP communities and SNMP users can be configured to use a view. When accessing SNMP through a community or users, access will be limited to OIDs included in the view. By default, communities and users can use default **efav3View** view of XCO.

About This Task

Follow this procedure to configure SNMP view and destination UDP port.

Table 22: Drift Reconcile & Idempotency Support

Identify Drift	Reconcile Configuration	Idempotency
Yes	Yes	Yes

Procedure

1. Create SNMP view.
   a. Run the following command to create an SNMP view:

```
efa inventory device snmp view create [ --ip device-ips | --name view-name | --mib-
tree mib-oid | --mib-tree-access access

--ip device-ips
 Comma separated range of device IP addresses. Example: 1.1.1.1-3,1.1.1.2,2.2.2.2
--name view-name
 View name
--mib-tree mib-oid
 MIB subtree in the form of Object identifier. Example: 1.3.6.1
--mib-tree-access access
 Mib-tree access. Valid values are: included, excluded
```

   The following example creates a view on a specified device:

```
efa inventory device snmp view create --ip 10.139.44.153-154 --name view1 --mib-
tree 1.3.6.1 --mib-tree-access included
+---------------+-------+----------+-----------------+---------+--------+
|  IP Address   | Name  | MIB-Tree | MIB-Tree-Access | Status  | Reason |
+---------------+-------+----------+-----------------+---------+--------+
| 10.139.44.153 | view1 | 1.3.6.1  | included        | Success |        |
+---------------+-------+----------+-----------------+---------+--------+
| 10.139.44.154 | view1 | 1.3.6.1  | included        | Success |        |
+---------------+-------+----------+-----------------+---------+--------+

Snmp view details
```

   b. Run the following command to delete an SNMP view:

```
efa inventory device snmp view delete [ --ip device-ips | --name view-name | --mib-
tree mib-oid |

--ip device-ips
 Comma separated range of device IP addresses. Example: 1.1.1.1-3,1.1.1.2,2.2.2.2
--name view-name
 View name
--mib-tree mib-oid
 MIB subtree in the form of Object identifier. Example: 1.3.6.1
```

   The following example deletes a view on a specified device:

```
efa inventory device snmp view delete --ip 10.139.44.153-154 --name view1 --mib-
tree 1.3.6.1
+---------------+-------+----------+---------+--------+
|  IP Address   | Name  | MIB-Tree | Status  | Reason |
+---------------+-------+----------+---------+--------+
| 10.139.44.153 | view1 | 1.3.6.1  | Success |        |
+---------------+-------+----------+---------+--------+
| 10.139.44.154 | view1 | 1.3.6.1  | Success |        |
+---------------+-------+----------+---------+--------+
```

```
Snmp view details
```

c. Run the following command to list SNMP view:

```
efa inventory device snmp view list [ --ip device-ips |

--ip device-ips
 Comma separated range of device IP addresses. Example: 1.1.1.1-3,1.1.1.2,2.2.2.2
```

The following example shows the current SNMP view for the specified device:

```
efa inventory device snmp view list --ip 10.139.44.153-154
+---------------+-------+----------+-----------------+----------+
|  IP Address   | Name  | MIB-Tree | MIB-Tree_Access | AppState |
+---------------+-------+----------+-----------------+----------+
| 10.139.44.153 | view1 | 1.2.3.4  | included        |          |
+               +-------+----------+-----------------+----------+
|               | view2 | 1.2.3.4  | included        |          |
+---------------+-------+----------+-----------------+----------+
| 10.139.44.154 | view3 | 1.2.3.5  | excluded        |          |
+---------------+-------+----------+-----------------+----------+

Snmp view details
```

d. Run the following command to create an SNMP community and SNMP group:

```
efa inventory device snmp community create [ --ip device-ips | --name community |
--group group | --enable-read-access | --enable-write
access | --enable-notify-access | --view view-name ]


--ip device-ip
Specifies a comma-separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2.
--name community
Specifies an SNMP community name.
--group group
Specifies an SNMP group name.
--enable-read-access
Sets read access for the view.
--enable-write-access
Sets write access for the view.
--enable-notify-access
Sets notify access for the view.
--view
Optionally specify a SNMP view name. Default view efav3View.used when not specified.
```

The following example creates a community using a specified device:

```
efa inventory device snmp community create --ip 10.139.44.153 --name community1 --
group group1 --enable-read-access --view view1
+---------+----------------+-------+------+------+--------+------+--------+-------+
|  IP     |    Community   |Group  |Read  |Write | Notify |View  |Status  |Reason |
|  Address|    Name        |       |view  |view  | view   |      |        |       |
+---------+----------------+-------+------+------+--------+------+--------+-------+
| 10.139  | $9$smklvisSghO |group1 |view1 |      |        |view1 |Success |       |
| .44.153 | ZEQvXJKBDeA==  |       |      |      |        |      |        |       |
+---------+----------------+-------+------+------+--------+------+--------+-------+
Snmp community details
```

e. Run the following command to list an SNMP community:

```
efa inventory device snmp community list [--ip device-ip ]

--ip device-ip
```

```
Specifies a comma-separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2
```

The following example creates a community using a specified device:

```
efa inventory device snmp community list --ip 10.139.44.153
+---------+----------------+-------+-------+------+--------+------+------------+
| IP      |    Community   |Group  | Read  |Write | Notify |View  | AppState   |
| Address |    Name        |       | view  |view  | view   |      |            |
+---------+----------------+-------+-------+------+--------+------+------------+
| 10.139  | $9$smklvisSghO |group1 | view1 |      |        |view1 | cfg-in-sync|
| .44.153 | ZEQvXJKBDeA==  |       |       |      |        |      |            |
+---------+----------------+-------+-------+------+--------+------+------------+
```

f. Run the following command to create an SNMP user and SNMP group:

```
efa inventory device snmp user create [--ip device-ip | --name community
| --group group | --enable-read-access | --enable-write-access | --
enable-notify-access | --auth-protocol md5 | sha | --auth-pass
authphrase | --priv-protocol AES128 | DES | --priv-pass privphrase |
--view view-name]

Parameters
--ip device-ip
Specifies a comma-separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2.
--name community
Specifies an SNMP community name.
--group group
Specifies an SNMP group name.
--enable-read-access
Sets read access for the view.
--enable-write-access
Sets write access for the view.
--enable-notify-access
Sets notify access for the view.
--auth-protocol md5 | sha
Sets notify access for the view. This parameter is set to off, by default.
--auth-pass passphrase
Authentication password.
--priv-protocol AES128 | DES
Privacy protocol.
--priv-pass privphrase
Privacy password.
--view view-name
Optionally specify a SNMP view name. Default view efav3View.used when not specified.
```

The following example creates users using a specified device:

```
efa inventory device snmp user create --ip 10.139.44.153 --name user1 --group
group1 --view view1
+--------+-----+------+-----+------+-------+-----+------+-----------+------
+-----------+--------+-------+
|IP      |User |Group |Read |Write |Notify |View |Auth
|Auth       |Priv  |Priv       |Status  |Reason |
|Address |     |      |view |view  |view   |     |proto
|passphrase |proto |passphrase |        |       |
+--------+-----+------+-----+------+-------+-----+------+-----------+------
+-----------+--------+-------+
|10.139. |user1|group1|     |      |       |View1|      |
|      |           |Success |       |
|44.153  |     |      |     |      |       |     |
|           |      |           |        |       |
+--------+-----+------+-----+------+-------+-----+------+-----------+------
```

```
+-----------+--------+-------+
Snmp user details
```

g. Run the following command to list an SNMP users:

```
efa inventory device snmp user list [--ip device-ip ]

--ip device-ip
Specifies a comma-separated range of device IP addresses. Example:
1.1.1.1-3,1.1.1.2,2.2.2.2
```

The following example list SNMP users:

```
efa inventory device snmp user list --ip 10.139.44.153
+--------+-----+------+-----+------+-------+------+-----------+------+-----------
+-----+------------+
|IP      |User |Group |Read |Write |Notify |View  |Auth       |Auth  |Priv
|Priv |AppState    |
|Address |     |view  |view |view  |       |proto |passphrase |proto |passphrase
|     |            |
+--------+-----+------+-----+------+-------+------+-----------+------+-----------
+-----+------------+
|10.139  |user1|group1|     |      |       |View1 |           |      |
|     |cfg-in-sync |
|.44.153 |     |      |     |      |       |      |           |      |
|     |            |
+--------+-----+------+-----+------+-------+------+-----------+------+-----------
+-----+------------+
Snmp user details
```

2. Create SNMP host.

a. Run the following command to create SNMP v2c or v3 host with a specified UDP port:

```
efa inventory device snmp host create [--ip device-ip | --host-ip IPv4 | IPv6 |
FQDN |--community community | --user user | --version v2c | v3
| --notify-type traps | informs | --engine-id remote_id | --udp-port port]

--ip device-ip                    Specifies a comma-separated range of device IP
                                  addresses. Example: 1.1.1.1-3,1.1.1.2,2.2.2.2.
--host-ip IPv4 | IPv6 | FQDN     Specifies a host IP address.
--community community             Specifies a community name. Applicable for v2c
                                  only.
--user user                       Specifies an SNMP v3 user.
--version v2c | v3                Specifies the SNMP version.
--notify-type traps | informs     Specifies the notification type. Informs are
                                  valid for v3 only.
--engine-id remote_id             Specifies the remote engine ID of manager.
--udp port                        Optional port number used to send notifications.
                                  Range: 0-65535, Default=162 (default 162)
```

The following example creates and lists SNMP host:

```
efa inventory device snmp host create --ip 10.139.44.153 --host-ip 1.1.1.1 --user
user1 --version v3 --notify-type traps --udp-port 163
+--------+-----+-----+----------+-------+-------+----------+---------+----+---------
+--------+-------+
|IP      |Host |User |Community |Notify |Engine |Source    | Vrf     |UDP |Severity
|Status  |Reason |
|Address |IP   |     |          |Type   |ID     |Interface |         |port|
|        |       |
+--------+-----+-----+----------+-------+-------+----------+---------+----+---------
+--------+-------+
|10.139  |1.1  |user1|          |traps  |       |management|mgmt-vrf |163 |None
|Success |       |
```

```
|.44.153 |.1.1 |     |           |       |       |chassis-ip|         |    |
|        |       |
+--------+-----+-----+----------+-------+-------+----------+---------+----+---------
+--------+-------+
Snmp host details

efa inventory device snmp host list --ip 10.139.44.153
+--------+-----+-----+----------+-------+---------+----------+---------+-----
+---------+------------+
|IP      |Host |User |Community |Notify |Remote   |Source    |
Vrf    |UDP  |Severity |AppState    |
|Address |IP   |     |          |Type   |EngineID |Interface |
|port |         |            |
+--------+-----+-----+----------+-------+---------+----------+---------+-----
+---------+------------+
|10.139  |1.1  |user1|          |traps  |         |management|mgmt-
vrf |163  |None     |cfg-in-sync |
|.44.153 |.1.1 |     |          |       |         |chassis-ip|
|163  |         |            |
+--------+-----+-----+----------+-------+---------+----------+---------+-----
+---------+------------+
Snmp host details
```

Drift and Reconcile (DRC) and Idempotency for SNMP

The table below captures the various attributes of the SNMP configuration interface for which DRC and idempotency is supported. A drift is identified if any of the fields below is modified through the SLX CLI or other management tool. A reconcile operation pushes the intended configuration to SLX, so keeping the SLX configuration in sync with XCO.

Regarding idempotency for creating an entry which already exists in XCO, an error message is returned stating that the user already exists.

Field	Identify Drift	Reconcile config	Idempotency	Comments
Community deleted	Yes	Yes	No	A valid error message is shown when a non-existent community is deleted.
Group name associated with community is modified	Yes	Yes	Not Applicable	
Group deleted	Yes	Yes	Not Applicable	
Modify group version.	No	No	Not Applicable	SLX does not support editing the SNMP group version.

Field	Identify Drift	Reconcile config	Idempotency	Comments
Modify read, review, or write view or notify view associated with group.	No	No	Not Applicable	SLX does not support editing the SNMP views associated with the group.
Modify groupname associated with SNMP user.	Yes	Yes	Not Applicable	
Modify authentication protocol associated with SNMP user.	Yes	Yes	Not Applicable	
Modify authentication password associated with SNMP user.	Yes	Yes	Not Applicable	
Modify privacy protocol associated with SNMP user.	Yes	Yes	Not Applicable	
Modify privacy password associated with SNMP user.	Yes	Yes	Not Applicable	
Delete SNMP user.	Yes	Yes	Not Applicable	A valid error message is shown when a non existent user is deleted.
Modify encrypted keyword associated with SNMP user.	Yes	Yes	Not Applicable	
Modify authentication type associated with group, meaning: auth, noauth, notify.	Yes	Yes	Not Applicable	
Delete SNMP host entry.	Yes	Yes	No	A valid error message is shown when a non existent host is deleted.

Field	Identify Drift	Reconcile config	Idempotency	Comments
Modify encrypted keyword associated with SNMP user.	Yes	Yes	Not Applicable	
Modify authentication type associated with group, meaning: auth, noauth, notify.	Yes	Yes	Not Applicable	
Delete SNMP host entry.	Yes	Yes	No	A valid error message is shown when a non existent host is deleted.
Update SNMP host security level.	No	No	Not Applicable	
Update SNMP host source interface.	No	No	Not Applicable	
Update SNMP host UDP port.	No	No	Not Applicable	
Update SNMP host VRF.	No	No	Not Applicable	
Update SNMP host engine id.	Yes	Yes	Not Applicable	
Update of SNMP host notification type [traps, informs]	Yes	Yes	Not Applicable	
Update of SNMP view MIB OID access [included, excluded]	Yes	Yes	Yes	
Delete SNMP view	Yes	Yes	Yes	

Unified Health and Fault Management

Unified Health and Fault Management Overview on page 503

Unified Health and Fault Management Overview

The Fault Management system raises alerts and alarms and maintains historical data. The Health Management system maintains the current health status of managed entities in XCO.

A Resource Path is used to identify an entity under the Health Management system and associate an alert raised by the Fault Management system.

XCO Unified View

The XCO Unified view allows the users to get the overall health of XCO and if unhealthy, identify the managed entity and infer the associated alert which has contributed to the unhealthy state. The XCO Unified View addresses all the managed entities as Resource Paths and represents the health status in a uniform response. This allows users to query and parse the health status quickly and consistently. All Fault Management and Health Management features are available for TPVM and Server deployments.

Hierarchical Representation of Resources

Top Level Hierarchy	Description
System	All the common software modules that make up the XCO infrastructure are covered under this node. For example, Database, K3s, Services, HA, and Certificate. There can be more sub-nodes under this as deemed necessary.
Component	This node groups all components (features) that are provided by XCO. For example, Visibility, Fabric, Tenant, Asset, and Common. Common groups all features like Firmware, License, and Fault. All entities are represented as resources in XCO so that they become addressable.

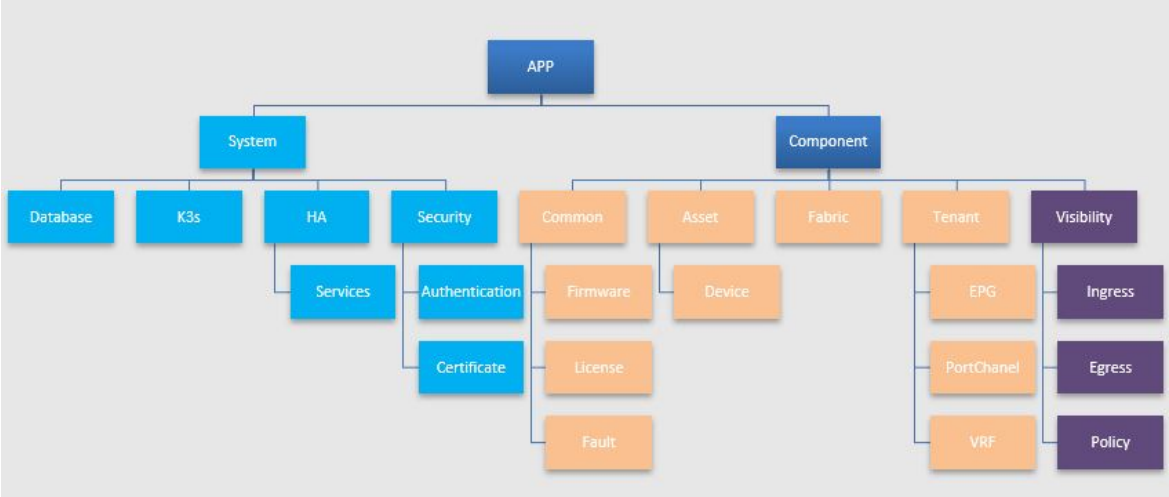

Figure 54: XCO Health Resource Hierarchy

Example:

- Certificate Management: `/App/System/Security/Certificate`
- Fabric named FabricOne: `/App/Component/Fabric?name=FabricOne`

Representing managed objects (for example, Fabric, Tenant, EPG, Devices, and Interfaces) and the software components (for example, License, Security, and Fault) as resource has the following advantages:

- Generic APIs can be provided for accessing the entities.
- All the resources can also be active participants in the system (for example, raise alerts or alarms).
- It provides a pluggable data model to support newer network types.
- It provides a deterministic way for modeling the operational model.

Unified View of Health and Fault Updates

Category of Events	Description
SLX Logs	• Device Logs coming from the device on RASLOG or GNMI Notification • Send as a Passthrough to Syslog over RELP or Webhook
Task Events	• Task level events generated by the XCO system • For example, Fabric Created and Tenant Modified
Status Updates	Health and Event updates from various components
Alerts	Alerts generated by Fault engine based on Status Updates from various components
Alarm	Stateful events generated by the fault rules engine

Note

SLX Logs and Task Events are delivered using the Notification Service.

API	Description
Event API	API for accessing Events
Health API	API for accessing Health

Fault Management - Alerts

Learn about all the possible alerts that are raised by Fault Management.

Common Alert Payload to be Published via Syslog

The following table provides the common fields of an alert object that are sent over the Syslog channel:

Field	SD-ID (Structured Data ID)	Example	Description	
<###>	N/A	116 =(14 * 8) + 4 Alert Range: 112-119	Priority Value: (Syslog Classifier * 8) + Syslog Severity	
			Syslog Classifier	
			14	log alert
			Syslog Severity	
			0	Emergency: system is unusable
			1	Alert: action must be taken immediately
			2	Critical: critical conditions
			3	Error: error conditions
			4	Warning: warning conditions
			5	Notice: normal but significant condition
			6	Informational: informational messages
			7	Debug: debug-level messages
Version	N/A	1	Version of syslog message	
Timestamp	N/A	2003-10-11T22:14:15.003Z	Timestamp of syslog message	
Hostname	N/A	xco.machine.com	Hostname of XCO	
App Name	N/A	faultmanager	Application generating syslog alerts	
Proc ID	N/A	-	Process ID	
Msg ID	N/A	-	Alert sub-type classification	
Sequence ID	meta	47	Tracks the sequence in which messages are submitted to the syslog transport.	
IP	origin	10.20.30.40	IP address (of XCO host)	

Field	SD-ID (Structured Data ID)	Example	Description	
Enterprise ID	origin	1916	Extreme Networks Enterprise ID	
Software	origin	XCO	Software Name (of XCO host)	
SW Version	origin	3.2.0	Software Version (of XCO host)	
Resource	alert@1916	/App/System/Security/Certificate?type=app_server_cert	XCO Health Resource path associated to the Alert being sent.	
Alert ID	alert@1916	31000	ID identifying the XCO Alert	
Cause	alert@1916	keyExpired	Reason for the Alert (Attempt to map to IANA standards)	
Type	alert@1916	securityServiceOrMechanismViolation	Indicates the Category (Attempt to map to IANA standards)	
Severity	alert@1916	warning	Severity of the XCO Alert (Critical, Major, Minor, Warning, Info)	
			XCO Alert	Syslog Severity
			Critical	Alert (1)
			Major	Critical (2)
			Minor	Error (3)
			Warning	Warning (4)
			Info	Informational (6)
BOMText	N/A	The application server certificate on the application will expire soon on “Sep 12 10:00:45 2022 GMT”.	(Byte Order Mask) Textual description of the Alert	

The following example maps alerts to RELP or Syslog fields:

```
<116>1 2003-10-11T22:14:15.003Z xco.machine.com faultmanager - -
   [meta sequenceId=”47”]
   [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]
   [alert@1916
    resource=”/App/System/Security/Certificate?type=app_server_cert”
    alertId=”31000”
    cause=”keyExpired”
    type=”securityServiceOrMechanismViolation”
    severity=”warning”]
   [alertData@1916
    type="app_server_cert"
    expiry_data="Sep 12 10:00:45 2022 GMT"]
   BOMThe application server certificate on the application will expire soon on “Sep 12
10:00:45 2022 GMT”.
```

Common Alert Payload to be published via Webhook

You can map alerts to Webhook payload.

```
{
  "type": "Alert",
  "timestamp": "2022-09-15T10:43:54.131202268-07:00",
  "severity": "major",
  "message": "Authentication failed for user \"root\".",
  "application": "faultmanager",
  "source_ip": "10.1.1.1",
  "device_ip": "",
  "username": "",
  "message_id": "",
  "hostname": "tpvm1",
  "logtype": "",
  "task": "",
  "scope": "",
  "status": "",
  "sequence_id": 7,
  "alert_id": 31010,
  "alarm_id": 0,
  "resource": "/App/System/Security/Authentication",
  "alarm_type": "securityOrMechanicalViolation",
  "alarm_cause": "credentialError",
  "alert_data": {
     "username": "root"
  }
}
```

Alert Commands

You can use alert commands to verify historical alerts and then filter the alerts based on resource, ID, severity or sequence ID.

The following table lists alert commands.

Command	Description
`efa system alert inventory show`	Alert inventory
`efa system alert inventory show --id=31000`	Alert definition filtered on Alert ID
`efa system alert inventory show --resource=/App/System/Certificate`	Alert definition filtered on Resource
`efa system alert show`	List of raised alerts
`efa system alert show {--id=31000 | --severity=warning | --sequence-id | --resource /App/System/Security/Certificate | --limit 5 | --before-timestamp | --after-timestamp }`	Alerts that are filtered by id or severity, sequence-id, resource or timestamps
`efa system alert show --before-timestamp 2022-08-11T17:32:28 --after-timestamp 2022-07-21T17:32:28`	List of alerts raised between some timestamps
`efa system alert show –limit=<number>`	Alerts from the alert history to the specified limit

Command	Description
`efa system alert show --id=31000`	Alerts from the history filtered on the ID
`efa system alert clear`	Clear the alert history

Inventory of Alerts

Use the information in the following tables to learn about the inventory of all possible alerts that are raised by Fault Management.

Note

- Fault Management does not generate alerts for any backup operations that are done for upgrade or restore.
- The `SystemBackupRestoreInitiatedAlert` is temporary and is visible only at the start of the restore process. This alert will not appear in the alert history after the restore is completed.

Alarm Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
MaximumAlarm InstancesReached	31900	Major	N/A - Alert does not affect health		

Backup and Restore Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
SystemBackupInitiatedAlert	31070	Info	N/A - Alert does not affect health		
SystemBackupSuccessAlert	31071	Info	N/A - Alert does not affect health		
SystemBackupFailureAlert	31072	Major	N/A - Alert does not affect health		
SystemBackupRestoreInitiatedAlert	31075	Info	N/A - Alert does not affect health		
SystemBackupRestoreSuccessAlert	31076	Info	N/A - Alert does not affect health		

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
SystemBackupRestoreFailureAler	31077	Major	N/A - Alert does not affect health		
FaultAlertSeqNumOutOfOrderAlert	31078	Info	N/A - Alert does not affect health		Informs that the sequence numbers might go out of order before this alert. The alert will be in order after the seqence number of this alert. This happens when backup is restored.

Certificate Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
CertificateExpiryNoticeAlert	31000	Warning	/App/System/Security/Certificate	type	app_server_cert
					default_intermediate_ca
					default_root_ca
					third_party_ca
					k3s_server_cert
					k3s_ca
					jwt_cert
DeviceCertificateExpiryNoticeAlert	31001	Warning	/App/System/Security/Certificate	device_ip	IP of the affected Device
				type	https_server_cert
					syslog_ca
					jwt_verifier
CertificateExpiredAlert	31002	Critical	/App/System/Security/Certificate	type	app_server_cert
					default_intermediate_ca
					default_root_ca
					third_party_ca
DeviceCertificateExpiredAlert	31003	Critical	/App/System/Security/Certificate	device_ip	IP of the affected Device
				type	syslog_ca
					jwt_verifier

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
CertificateRenewalAlert	31004	Info	/App/System/Security/Certificate	type	app_server_cert
					default_intermediate_ca
					default_root_ca
					third_party_ca
					jwt_cert
DeviceCertificateRenewalAlert	31005	Info	/App/System/Security/Certificate	device_ip	IP of the device
				type	https_server_cert
					jwt_verifier
CertificateUnreadableAlert	31006	Warning	/App/System/Security/Certificate	type	app_server_cert
					default_intermediate_ca
					default_root_ca
					third_party_ca
					k3s_server_cert
					k3s_ca
					jwt_cert
DeviceCertificateUnreadableAlert	31007	Warning	/App/System/Security/Certificate	device_ip	IP of the device type
				type	https_server_cert
					syslog_ca
					jwt_verifier
DeviceCertificateDeviceRemovedAlert	31008	Info	/App/System/Security/Certificate	device_ip	IP of the affected Device
				type	https_server_cert
					syslog_ca
					jwt_verifier

Device Connectivity Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
DeviceConnectivityFailureAlert	31501	Major	/App/Component/Asset/Device	device_ip	IP of the device
DeviceConnectivitySuccessAlert	31502	Info	/App/Component/Asset/Device	device_ip	IP of the device
DeviceConnectivityDeviceRemovedAlert	31503	Info	/App/Component/Asset/Device	device_ip	IP of the device

High Availability Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
HAServiceNonRedundantAlert	31050	Major	/App/System/HA/Nodes/Node		
HAServiceFullyRedundantAlert	31051	Info	/App/System/HA/Nodes/Node		
HAServiceNewActiveAlert	31052	Info	/App/System/HA/Nodes/Node		
ServiceDegradedAlert	31053	Warning	/App/System/HA/Nodes/Services		
ServiceRestoredAlert	31054	Info	App/System/HA/Nodes/Services		

Table 23: Login Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
LoginFailureAlert	31010	Major	/App/System/Security/Authentication (Alert does not affect health)		
LoginSuccessfulAlert	31011	Info	/App/System/Security/Authentication (Alert does not affect health)		

LDAP Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
LDAPServerConnectivityFailureAlert	31030	Major	/App/System/Security/Authentication	server	LDAP Server IP address
LDAPServerConnectivitySuccessAlert	31031	Info	/App/System/Security/Authentication	server	LDAP Server IP address
LDAPServerConfigurationRemovedAlert	31033	info	/App/System/Security/Authentication	server	LDAP Server IP address

Storage Utilization Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
StorageUtilizationThresholdAlert	31040	Warning	/App/System/Storage	node_ip	IP address of the node
				mount_point	Name of the mount point
StorageUtilizationFullAlert	31041	Critical	/App/System/Storage	node_ip	IP address of the node
				mount_point	Name of the mount point
StorageUtilizationCheckAlert	31042	Info	/App/System/Storage	node_ip	IP address of the node
				mount_point	Name of the mount point

Upgrade Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
UpgradeAsyncOperationInitiatedAlert	31060	Info	N/A - Alert does not affect health		
UpgradeAsyncOperationSuccessAlert	31061	Info	N/A - Alert does not affect health		

Alert Details

The following topics describe all possible alerts in detail that are raised by Fault Management.

- Backup and Restore Alerts on page 514
- Certificate Alerts on page 521
- Device Connectivity Alerts on page 529
- High Availability Alerts on page 531
- Login Alerts on page 536
- LDAP Alerts on page 538
- Storage Alerts on page 542
- Upgrade Alerts on page 546

Alarm Alerts

Use the information in the following tables to learn about all possible login alerts in detail that are raised by Fault Management.

Maximum Alarm Instances Reached

31900	Maximum Alarm Instances Reached
Description	Send an alert when the number of alarms created reach the maximum capacity.
Preconditions	On TPVM, the maximum number of alarm instances that can exist is equal to one less than its maximum capacity.

Requirements	Alert shows the following data:
	• Deployment Platform
	• Max Instances
	The following example shows an alert when the device reaches maximum number of alarms:
	`<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]`
	`[alert@1916`
	`resource=””`
	`alertId=”31900”`
	`cause=”resourceAtOrNearingCapacity”`
	`type=”processingErrorAlarm”`
	`severity=”major”]`
	`[alertData@1916`
	`deployment_platform=”TPVM”`
	`max_instances=”500”]`
	`BOMThe application has reached the maximum of “500" alarms on the “TPVM” platform. Close and purge alarms to free up resources.`

Health Response	Health resources are not associated with max alarm instances reached.

Backup and Restore Alerts

Use the information in the following tables to learn about all possible backup and restore alerts that are raised by Fault Management.

System Backup Initiated

31070	System Backup Initiated
Description	Send an alert when a backup is initiated.
Preconditions	None

31070	System Backup Initiated
Requirements	Alert shows the following data: • User • Version • Backup Type • Fabric Name • Fabric All • Device IPs The following example shows an alert when a system backup is initiated: `<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”” alertId=”31070” cause=”operationInitiated” type=”communicationsAlarm” severity=”info”] [alertData@1916 user=”extreme” version=”3.2.0-backup” backup_type=”FabricName” fabric_name=”nonclos” fabric_all=”” device_ips=””] BOMSystem backup type "FabricName" initiated for fabric "nonclos" for efa version "3.2.0-backup" by user "extreme".`
Health Response	Health resources are not associated with system backup.

System Backup Success

31071	System Backup Success
Description	Send an alert when backup is successful.
Preconditions	You initiated a system backup or a periodic backup.

31071	System Backup Success
Requirements	Alert shows the following data: • User • Version • Backup Type • Fabric Name • Fabric All • Filename The following example shows an alert when a system backup is successful:: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”” alertId=”31071” cause=”operationSuccess” type=”communicationsAlarm” severity=”info”] [alertData@1916 user=”extreme” version=”3.2.0-backup” backup_type=”FabricAll” fabric_name=”” fabric_all=true device_ips=”” filename="/apps/efa_logs/backup/XCO-3.2.0-backup-2022-12-07T19-05-31.719.tar”] BOMSystem backup type "FabricAll" completed successfully for efa version "3.2.0-backup" by user "extreme" located at "/apps/efa_logs/backup/XCO-3.2.0-backup-2022-12-07T19-05-31.719.tar".
Health Response	Health resources are not associated with system backup.

System Backup Failure

31072	System Backup Failure
Description	Send an alert when the backup operation fails.
Preconditions	You initiated a system backup or a periodic backup.

31072	System Backup Failure

Requirements	Alert shows the following data:
	• User
	• Version
	• Backup Type
	• Fabric Name
	• Fabric All
	• Device IPs
	• Error
	The following example shows an alert when backup fails for a device:
	`<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]`
	`[alert@1916`
	`resource=””`
	`alertId=”31072”`
	`cause=”operationFailure”`
	`type=”processingErrorAlarm”`
	`severity=”major”]`
	`[alertData@1916`
	`user=”extreme”`
	`version=”3.2.0-backup”`
	`backup_type=”DeviceIP”`
	`fabric_name=””`
	`fabric_all=””`
	`device_ips=”10.24.80.55,10.24.80.57”`
	`error=” Config backup Failed for device: [10.24.80.55,10.24.80.57], Reason: Devices do not exist”]`
	`BOMSystem backup type "DeviceIP" failed for device ips "10.24.80.55,10.24.80.57" for efa version "3.2.0-backup" by user "extreme" due to " Config backup Failed for device: [10.24.80.55,20.24.80.57], Reason: Devices do not exist ".`

Health Response	Health resources are not associated with system backup.

System Backup Restore Initiated

31075	System Backup Restore Initiated
Description	Send an alert when the backup restore operation is initiated.
Preconditions	None

31075	System Backup Restore Initiated
Requirements	Alert shows the following data: • User • Version • Filename The following example shows an alert when a system restore is initiated from a backup: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”” alertId=”31075” cause=”operationInitiated” type=”communicationsAlarm” severity=”info”] [alertData@1916 user=”extreme” version=”3.2.0-backup” filename=” XCO-3.2.0-backup-2022-12-05T15-48-54.920.tar”] BOMSystem restore initiated with "XCO-3.2.0-backup-2022-12-05T15-48-54.920.tar" on XCO version "3.2.0-backup" by user "extreme".
Health Response	Health resources are not associated with system restore.

System Backup Restore Success

31076	System Backup Restore Success
Description	Send an alert when the restore operation is successful
Preconditions	You initiated a system restore.

31076	System Backup Restore Success

Requirements	Alert shows the following data:
	• User: User who initiated this restore operation
	• Version: version
	• Filename: Backup file that was restored
	The following example shows an alert when a system restore is successful:
	`Syslog RFC-5424 Example:`
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]`
	`[alert@1916`
	`resource=""`
	`alertId="31076"`
	`cause="operationSuccess"`
	`type="communicationsAlarm"`
	`severity="info"]`
	`[alertData@1916`
	`user="extreme"`
	`version="3.2.0-backup"`
	`filename=" XCO-3.2.0-backup-2022-12-05T15-48-54.920.tar"]`
	`BOMSystem restore completed successfully with "XCO-3.2.0-backup-2022-12-05T15-48-54.920.tar" on XCO version "3.2.0-backup" by user "extreme"`

Health Response	Health resources are not associated with system restore.

System Backup Restore Failure

31077	System Backup Restore Failure
Description	Send an alert when a system restore fails.
Preconditions	You initiated a system restore.

31077	System Backup Restore Failure
Requirements	Alert shows the following data: • User • Version • Filename • Error The following example shows an alert when the system retire fails from a backup: <114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="" alertId="31077" cause="operationFailure" type="processingErrorAlarm" severity="major"] [alertData@1916 user=extreme"" version="3.2.0-backup" filename=" XCO-3.2.0-backup-2022-12-05T15-48-54.920.tar" error="reason for failure"] BOMSystem restore failed with "XCO-3.2.0-backup-2022-12-05T15-48-54.920.tar" on XCO version "3.2.0-backup" by user "extreme"
Health Response	Health resources are not associated with system restore.

System Backup Restore Out of Order Sequence ID

31078	System Backup Restore Out of Order Sequence ID
Description	Send an alert when the sequence number is out of order due to system backup and restore. The alert indicates that the existing sequence number is out of order, and a new sequence number from the backup tar file takes effect.
Preconditions	Backup tar file already has the alerts.

31078	System Backup Restore Out of Order Sequence ID
Requirements	Alert shows the following data: • None The following example shows an alert when the sequence number goes out of order for backup and restore: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="" alertId="31078" cause="informationOutOfSequence" type="integrityViolation" severity="info"] BOMAlert sequence number can go out of order.
Health Response	Health resources are not associated with out of order system sequence.

Certificate Alerts

Use the information in the following tables to learn about all possible certificate alerts in detail that are raised by Fault Management.

XCO Certificate Expiry Notice

31000	XCO Certificate Expiry Notice
Description	Send an alert when an XCO certificate is about to expire.
Preconditions	You cannot configure the system default settings in Certificate Manager component. • Polling frequency for certificate expiry notice is **daily**. • Monitors the following types of XCO certificate and its value: ◦ App Server Certificate (of XCO): app_server_cert ◦ Default Intermediate CA: default_intermediate_ca ◦ Default Root CA: default_root_ca ◦ Third-Party CA: third_party_ca ◦ K3s Server Certificate: k3s_server_cert ◦ K3s CA: k3s_ca ◦ JWT Certificate: jwt_cert The polling service sends the "CertificateExpiryNoticeAlert" notification with an expiry date.

31000	XCO Certificate Expiry Notice
Requirements	Alert shows the following data: • Certificate Type • Expiry Date The following example shows an alert when an XCO certificate (for example, App Server Certificate) is about to expire: `<116>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/Security/Certificate?type=app_server_cert" alertId="31000" cause="keyExpired" type="securityServiceOrMechanismViolation" severity="warning"] [alertData@1916 type="app_server_cert" expiry_date="Sep 12 10:00:45 2022 GMT"] BOMThe App Server Certificate on the application will expire soon on "Sep 12 10:00:45 2022 GMT".`
Health Response	Response `{ Resource: /App/System/Security/Certificate?type=app_server_cert HQI { Color: Yellow Value: 1 } StatusText: The App Server Certificate on the application will expire soon on "Sep 12 10:00:45 2022 GMT". }`

Managed Device Certificate Expiry Notice

31001	Managed Device Certificate Expiry Notice
Description	Send an alert when a certificate on the SLX device is about to expire.
Preconditions	You cannot configure the default system settings in Inventory Service. • Polling frequency for certificate expiry notice is **daily** • Monitors the following types of Device Certificate and its value: ◦ HTTPS Server Certificate: https_server_cert ◦ Syslog CA: syslog_ca ◦ JWT Verifier (OAuth2): jwt_verifier The polling service sends the "DeviceCertificateExpiryNoticeAlert" notification with an expiry date.

31001	Managed Device Certificate Expiry Notice

Requirements	Alert shows the following data:
	• Device IP
	• Certificate Type
	• Expiry Date
	
	The following example shows an alert when a certificate (for example, HTTPS Server Certificate) is about to expire on SLX device:
	`<116>1 2022-10-11T22:14:15.003Z xco.machine.com FaultManager -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_cert"`
	`alertId="31001"`
	`cause="keyExpired"`
	`type="securityServiceOrMechanismViolation"`
	`severity="warning"]`
	`[alertData@1916`
	`device_ip="10.10.10.1"`
	`type="https_server_cert"`
	`expiry_date="Sep 12 10:00:45 2022 GMT"]`
	`BOMThe HTTPS Server Certificate on device "10.10.10.1" will expire soon on "Sep 12 10:00:45 2022 GMT".`

Health Response	Response
	`{`
	`Resource:/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_cert`
	`HQI {`
	`Color: Yellow`
	`Value: 1`
	`}`
	`StatusText: The HTTPS Server Certificate on device "10.10.10.1" will expire soon on "Sep 12 10:00:45 2022 GMT".`
	`}`

XCO Certificate Expired

31002	XCO Certificate Expired
Description	Send an alert when an XCO certificate has expired. You will not get this alert when the system is not functional.
Preconditions	K3s must be up and running Only supports non-k3s cert expiry. • Polling frequency for certificate expiry notice is **daily** • Monitors the following types of XCO Certificate and its value: ◦ App Server Certificate (of XCO): app_server_cert ◦ Default Intermediate CA: default_intermediate_ca ◦ Default Root CA: default_root_ca ◦ Third-Party CA: third_party_ca When the App Server Certificate expires, you cannot communicate with XCO via REST API. Therefore, you cannot query the health status.

Requirements	Alert shows the following data:
	• Certificate Type
	• Expired Date
	
	The following example shows an alert when an XCO certificate (for example, App Server Certificate) is expired:
	`<113>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?type=app_server_cert"`
	`alertId="31002"`
	`cause="keyExpired"`
	`type="securityServiceOrMechanismViolation"`
	`severity="critical"]`
	`[alertData@1916`
	`type="app_server_cert"`
	`expire_date="Sep 12 10:00:45 2022 GMT"]`
	`BOMThe App Server Certificate on the application has expired on "Sep 12 10:00:45 2022 GMT".`

Health Response	Response
	`{`
	`Resource: /App/System/Security/Certificate?type=app_server_cert`
	`HQI {`
	`Color: Black`
	`Value: 4`
	`}`
	`StatusText: The App Server Certificate on the application has expired on "Sep 12 10:00:45 2022 GMT".`
	`}`

Managed Device Certificate Expired

31003	Managed Device Certificate Expired
Description	Send an alert when an SLX certificate has expired
Preconditions	To allow the RASLog service to receive events from an SLX device, ensure the device is registered and the SLX syslog server configuration points to the XCO IP. When a syslog CA certificate expires, SLX device does not send the syslog alerts to the RASLog service. • Polling frequency for certificate expiry notice is **daily**. • Monitors the following types of Device Certificate and its value: ◦ Syslog CA: syslog_ca ◦ JWT Verifier (OAuth2): jwt_verifier The polling service sends the “DeviceCertificateExpiredNoticeAlert” notification with an expiry date.

Requirements	Alert shows the following data:
	• Device IP
	• Certificate Type
	• Expired Date
	The following example shows an alert when an SLX certificate (for example, Syslog CA) is expired:
	`<113>1 2022-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]`
	`[alert@1916`
	`resource=”/App/System/Security/Certificate?device_ip=10.10.10.1&type=syslog_ca”`
	`alertId=”31003”`
	`cause=”keyExpired”`
	`type=”securityServiceOrMechanismViolation”`
	`severity=”critical”]`
	`[alertData@1916`
	`device_ip=”10.10.10.1”`
	`type=”syslog_ca”`
	`expiry_date=”Sep 12 10:00:45 2022 GMT”]`
	`BOMThe Syslog CA on device “10.10.10.1” has expired on “Sep 12 10:00:45 2022 GMT”`

Health Response	*Response*
	`{`
	`Resource:/App/System/Security/Certificate?device_ip=10.10.10.1&type=syslog_ca`
	`HQI {`
	`Color: Black`
	`Value: 4`
	`}`
	`StatusText: The Syslog CA on device “10.10.10.1” has expired on “Sep 12 10:00:45 2022 GMT.`
	`}`

XCO Certificate Upload or Renewal

31004	XCO Certificate Upload or Renewal
Description	Send an alert when a certificate is renewed.
Preconditions	• Sends an alert for renewal of the certificates managed by XCO. • XCO sends a renewal alerts for the following types of certificate and its value: ◦ App Server Certificate (of XCO): app_server_cert ◦ Default Intermediate CA: default_intermediate_ca ◦ Default Root CA: default_root_ca ◦ Third-Party CA: third_party_ca ◦ JWT Certificate: jwt_cert ◦ K3s Server Certificate: k3s_server_cert ◦ K3s CA Certificate: k3s_ca

Requirements	Alert shows the following data:
	• Certificate Type
	The following example shows an alert when an XCO certificate is renewed:
	`Syslog RFC-5424 Example:`
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?type=app_server_cert"`
	`alertId="31004"`
	`cause="keyGenerated"`
	`type="securityServiceOrMechanismViolation"`
	`severity="warning"]`
	`[alertData@1916`
	`type="app_server_cert"]`
	`BOMThe App Server Certificate on the application has bee renewed.`

Health Response	Response
	`{`
	`Resource: /App/System/Security/Certificate?type=app_server_cert`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The App Server Certificate on the application has been renewed.`
	`}`

Managed Device Certificate Upload or Renewal

31005	Managed Device Certificate Upload or Renewal
Description	Send an alert when a device certificate is renewed.
Preconditions	Sent an alert on renewal of following certificates on devices: • HTTPS Server Certificate: https_server_cert • JWT Verifier (OAuth2): jwt_verifier

Requirements	Alert shows the following data:
	• Device IP
	• Certificate Type
	The following example shows an alert when a device certificate is renewed:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_cert"`
	`alertId="31005"`
	`cause="keyGenerated"`
	`type="securityServiceOrMechanismViolation"`
	`severity="info"]`
	`[alertData@1916`
	`device_iP="10.10.10.1"`
	`type="https_server_cert"]`
	`BOMThe HTTPS Server Certificate on the device 10.10.10.1 has been renewed.`

Health Response	Response
	`{`
	`Resource:/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_cert`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The HTTPS Server Certificate on the device 10.10.10.1 has been renewed.`
	`}`

Managed Device Certificate Expiration Device Removed

31008	Managed Device Certificate Expiration Device Removed
Description	Send an alert when an SLX device is removed from a managed device
Preconditions	The SLX device is registered in inventory service. • You can run a command for device removal from inventory service. • Monitors the following types of Device Certificates: ◦ HTTPS Server Certificate: https_server_cert ◦ Syslog CA: syslog_ca ◦ JWT Verifier (OAuth2): jwt_verifier The removed device sends three alerts to clear any unhealthy state in the health service.

Requirements	Alert shows the following data:
	• Device IP
	• Certificate Type
	The following example shows an alert when an SLX device is removed:
	`<118>1 2022-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_cert"`
	`alertId="31008"`
	`cause="configRemoved"`
	`type="securityServiceOrMechanismViolation"`
	`severity="info"]`
	`[alertData@1916`
	`device_ip="10.10.10.1"`
	`type="https_server_cert"]`
	`BOMThe device 10.10.10.1 has been removed so cleaning up HTTPS Server Certificate`

Health Response	Response
	`{`
	`Resource:/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_cert`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The device 10.10.10.1 has been removed so cleaning up HTTPS Server Certificate.`
	`}`

Device Connectivity Alerts

Use the information in the following tables to learn about all possible device connectivity alerts in detail that are raised by Fault Management.

Managed Device Connectivity Loss

31501	Managed Device Connectivity Loss
Description	Send an alert when XCO loses contact with SLX.
Preconditions	The device is registered and the connectivity between devices is verified during periodic device updates. The polling for connectivity occurs when the device health check is enabled. The following command output is an example of user configuration: `efa inventory device setting update --ip=10.10.10.1 --health-check-enable Yes --health-check-interval 6m --health-check-heartbeat-miss-threshold 5` The polling service sends the “DeviceConnectivityFailureAlert” notification when there is a loss of contact.
Requirements	Alert shows the following data: • Device IP • Failed Adapters • Failure Reason The following example shows an alert when XCO contact with SLX device is lost: `<114>1 2022-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”/App/Component/Asset/Device?device_ip=10.10.10.1” alertId=”31501” cause=”connectionEstablishmentError” type=”communicationsAlarm” severity=”major”] [alertData@1916 device_ip=”10.10.10.1” failed_adapters=”ssh rest netconf” failure_reason=”Authentication failed” BOMContact has been lost with device “10.10.10.1”`
Health Response	Response `{ Resource: /App/Component/Asset/Device?device_ip=10.10.10.1 HQI { Color: Red Value: 3 } StatusText: Contact has been lost with device "10.10.10.1". }`

Managed Device Connectivity Reestablished

31502	Managed Device Connectivity Reestablished
Description	Send an alert when the SLX device is reachable.
Preconditions	The device is registered, and the connectivity is checked during normal periodic device updates. The polling for connectivity occurs when the device health check is enabled. The following is a sample example of user configuration: `efa inventory device setting update --ip=10.10.10.1 --health-check-enable Yes --health-check-interval 6m --health-check-heartbeat-miss-threshold 5` The polling service sends the “DeviceConnectivitySuccessAlert” notification when an SLX device is not reachable.
Requirements	Alert shows the following data: • Device IP The following example shows an alert when a device is unreachable: `<118>1 2022-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”/App/Component/Asset/Device?device_ip=10.10.10.1” alertId=”31502” cause=”connectionEstablished” type=”communicationsAlarm” severity=”info”] [alertData@1916 device_ip=”10.10.10.1” BOMContact has been regained with device “10.10.10.1”.`
Health Response	Response `{ Resource: /App/Component/Asset/Device?device_ip=10.10.10.1 HQI { Color: Green Value: 0 } StatusText: Contact has bee regained with device "10.10.10.1". }`

Managed Device Connectivity Device Removed

31503	Managed Device Connectivity Device Removed
Description	Send an alert when the SLX device is removed.
Preconditions	The “DeviceConnectivityDeviceRemovedAlert” notification is sent when a device is removed.

31503	Managed Device Connectivity Device Removed
Requirements	Alert shows the following data: • Device IP The following example shows an alert when a device is removed: <118>1 2022-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/Component/Asset/Device?device_ip=10.10.10.1" alertId="31503" cause="configRemoved" type="communicationsAlarm" severity="info"] [alertData@1916 device_ip="10.10.10.1" BOMDevice "10.10.10.1" has been removed.

Health Response	Response
	{
	Resource: /App/Component/Asset/Device?device_ip=10.10.10.1
	HQI {
	Color: Green
	Value: 0
	}
	StatusText: Device "10.10.10.1" has been removed.
	}

High Availability Alerts

Use the information in the following tables to learn about all possible HA alerts in detail that are raised by Fault Management.

HA Service (Non-Redundant)

31050	HA Service (Non-Redundant)
Description	Send an alert when the standby node is not up which indicates that the system has no redundancy.
Preconditions	Starting with EFA 3.1.0, a timer task periodically monitors the status of the standby node, and raises an event to the fault management system. The fault management system raises an alert to the user to indicate that the system is not fully redundant. • For HA events, the polling frequency is **every minute**.

31050	HA Service (Non-Redundant)
Requirements	Alert shows the following data: • Node IP The following example shows an alert when the standby node is down: <114> 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/HA/Nodes/Node" alertId="31050" cause="lossOfRedundancy" type="operationalViolation" severity="major"] [alertData@1916 node_ip="10.1.2.4"] BOMHA degraded, node 10.1.2.4 is down.

Health Response	Response
	{
	Resource: /App/System/HA/Nodes/Node
	HQI {
	Color: Red
	Value: 3
	}
	StatusText: HA degraded, node 10.2.3.5 is down.
	}

HA Service (Fully Redundant)

31051	HA Service (Fully Redundant)
Description	Send an alert when the standby node is up and ready which indicates that the system is fully redundant.
Preconditions	A timer task periodically monitors the status of the nodes and raises an event to the fault management system. The fault management system raises an alert to the user to indicate that the system is fully redundant. • For HA events, the polling frequency is **every minute**.

31051	HA Service (Fully Redundant)
Requirements	Alert shows the following data: • None The following example shows an alert when the standby node is up and running: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/HA/Nodes/Node" alertId="31051" cause="redundancyRestored" type="operationalViolation" severity="info"] BOMHA fully redundant

Health Response	Response
	{
	Resource: /App/System/HA/Nodes/Node
	HQI {
	Color: Green
	Value: 0
	}
	StatusText: HA fully redundant.
	}

HA Service (Failover Occurred)

31052	HA Service (Failover Occurred)
Description	Send an alert when an HA failover has occurred.
Preconditions	A timer task periodically monitors the status of the nodes and raises an event to the fault management system. The fault management raises an alert to the user to indicate that an HA failover has occurred. • For HA events, polling frequency is **every minute**.

31052	HA Service (Failover Occurred)
Requirements	Alert shows the following data: • Active IP The following example shows an alert when there is a HA failure: <114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/HA/Nodes/Node" alertId="31052" cause="localNodeTransmissionError" type="operationalViolation" severity="major"] [alertData@1916 active_iP="10.1.2.3"] BOM10.1.2.3 is now the HA active node

Health Response	Response
	{
	Resource: /App/System/HA/Nodes/Node
	HQI {
	Color: Red
	Value: 3
	}
	StatusText: 10.1.2.3 is now the HA active ndoe.
	}

Service Degraded

31053	Service Degraded
Description	Send an alert when some of the node services are not operational.
Preconditions	A timer task periodically monitors the node status and raises an event to the fault management system. The fault management system raises an alert to the user to indicate that some of the node services are not running. • For service events, the polling frequency is **every minute**.

31053	Service Degraded
Requirements	Alert shows the following data: • None None The following example shows an alert when some of the node services are not running: <116>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/HA/Nodes/Services" alertId="31053" cause="serviceDegraded" type="operationalViolation" severity="warning"] BOMSome of the services are not operational.

Health Response	Response
	{
	Resource: /App/System/HA/Nodes/Services
	HQI {
	Color: Yellow
	Value: 1
	}
	StatusText: Some of the services are not operational.
	}

Service Restored

31054	Service Restored
Description	Send an alert when all the node services are operational.
Preconditions	A timer task raises an event to the fault management system. The fault management system raises an alert to indicate to the user that some of the node services are running. • For service events, the polling frequency is **every minute**.

31054	Service Restored
Requirements	Alert shows the following data: • None The following example shows an alert when all the node services are running: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/HA/Nodes/Services" alertId="31054" cause="serviceRestored" type="operationalViolation" severity="info"] BOMServices are in running state.

Health Response	Response
	{
	Resource: /App/System/HA/Nodes/Services
	HQI {
	Color: Green
	Value: 0
	}
	StatusText: Services are in running state.
	}

Login Alerts

Use the information in the following tables to learn about all possible login alerts in detail that are raised by Fault Management.

Security Level Thresholds (Login attempts)

31010	Security Level Thresholds (Login attempts)
Description	Send an alert when a user login attempt to XCO fails
Preconditions	None

31010	Security Level Thresholds (Login attempts)
Requirements	Alert shows the following data: • Username The following example shows an alert when a login attempt to XCO fails: <114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 alertId=”31010” cause=”credentialError” type=”securityServiceOrMechanismViolation” severity=”major”] [alertData@1916 username=”bob”] BOMAuthentication failed for user “bob”.
Health Response	N/A

Login Successful

31011	Login Successful
Description	Send an alert when a user login attempt to XCO is successful.
Preconditions	None
Requirements	Alert shows the following data: • Username The following example shows an alert when a login attempt to XCO is successful: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”/App/System/Security/Authentication” alertId=”31011” cause="loginSuccessful" type="other" severity=”info”] [alertData@1916 username=”bob”] BOMAuthentication successful for user "bob".
Health Response	N/A

LDAP Alerts

Use the information in the following tables to learn about all possible LDAP alerts in detail that are raised by Fault Management.

LDAP Connectivity Failure

31030	LDAP Connectivity Failure
Description	Send an alert when LDAP server in XCO is not reachable
Preconditions	The polling is enabled only if: 1. LDAP servers are configured in the system. 2. Authentication fallback preference is set to LDAP. • System monitors all the LDAP servers. • Polling frequency is **hourly**. • During polling, login is performed with the base user configured in the system.

31030	LDAP Connectivity Failure
Requirements	Alert shows the following data: • Server • Name • Error The following example shows an alert when an LDAP server is not reachable: <114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/Security/Authentication?server=10.1.2.3" alertId="31030" cause="underlyingResourceUnavailable" type="communicationsAlarm" severity="major"] [alertData@1916 server="10.1.2.3" name="ldap1" error="failed to connect to the LDAP server with the given credentials"] BOMLDAP connectivity check failed for the server '10.x.x.x' configured with the name 'ldap1' due to 'failed to connect to the LDAP server with the given credentials.

Health Response	Response
	{
	Resource: /App/System/Security/Authentication?server=10.1.2.3
	HQI {
	Color: Red
	Value: 3
	}
	StatusText: LDAP connectivity check failed for the server '10.1.2.3' configured with the name 'ldap1' due to 'failed to connect to the LDAP server with the given credentials.
	}

LDAP Server Connectivity Success

31031	LDAP Server Connectivity Success
Description	Send an alert when LDAP server configured in XCO is reachable.
Preconditions	The polling is enabled only if: 1. LDAP servers are configured in the system. 2. Authentication fallback preference is set to LDAP. • System monitors all the LDAP servers. • Polling frequency is **hourly**. • During polling, login is performed with the base user configured in the system.

31031	LDAP Server Connectivity Success
Requirements	Alert shows the following data: • Server • Name The following example shows an alert when an LDAP server is reachable: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId=”47”] [origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”] [alert@1916 resource=”/App/System/Security/Authentication?server=10.1.2.3” alertId=”31031” cause=”connectionEstablished” type=”communicationsAlarm” severity=”info”] [alertData@1916 server=”10.1.2.3” name=”ldap1”] BOMLDAP connectivity check success for the server ‘10.1.2.3’ configured with the name ‘ldap1’.
Health Response	Response { Resource: /App/System/Security/Authentication?server=10.1.2.3 HQI { Color: Green Value: 0 } StatusText: LDAP connectivity check success for the server ‘10.1.2.3’ configured with the name ‘ldap1’. }

LDAP Server Configuration Removed

31032	LDAP Server Configuration Removed
Description	Send an alert when LDAP server configuration is removed
Preconditions	The polling is enabled only if: 1. LDAP servers are configured in the system. 2. Authentication fallback preference is set to LDAP. • System monitors all the LDAP servers. • Polling frequency is **hourly**. • During polling, login is performed with the base user configured in the system.

31032	LDAP Server Configuration Removed
Requirements	Alert shows the following data: • Server • Name The following example shows an alert when an LDAP server is removed: <118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/Security/Authentication?server=10.1.2.3" alertId="31032" cause="configRemoved" type="communicationsAlarm" severity="info"] [alertData@1916 server="10.1.2.3" name="ldap1"] BOMLDAP 'ldap1' with server ip '10.1.2.3' removed

Health Response	Response
	{
	Resource: /App/System/Security/Authentication?server=10.1.2.3
	HQI {
	Color: Green
	Value: 0
	}
	StatusText: LDAP 'ldap1' with server ip '10.1.2.3' removed.
	}

Storage Alerts

Use the information in the following tables to learn about all possible storage alerts in detail that are raised by Fault Management.

Storage Utilization Threshold

31040	Storage Utilization Threshold
Description	Send an alert for each TPVM mount point when capacity reaches 75% utilization or more.
Preconditions	You cannot configure the default settings in the System Component (Monitor Service). • Polling frequency is **hourly** for the storage utilization threshold notice. • The following are storage utilization thresholds: ◦ **Under 75% - info (31042 is raised)** ◦ **75% - warning** ◦ **85% - minor** ◦ **95% - major** ◦ **97% - critical** • System monitors the TPVM storage utilization on the following mount points: ◦ **“/” (/dev/vda2)** ◦ **“/apps” (/dev/vdb1)** The polling service sends an alert with the TPVM storage utilization percentage.

31040	Storage Utilization Threshold
Requirements	Alert shows the following data: • Node IP • Mount Point • Used MB • Available MB • Utilization Percent The following example shows an alert when a node IP mount point reaches 75% of storage utilization: `<116>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/Storage?node_ip=10.2.3.4&mount_point=/" alertId="31040" cause="storageCapacityProblem" type="processingErrorAlarm" severity="warning"] [alertData@1916 node_ip="10.2.3.4" mount_point="/" used_mb="7114" available_mb="2371" utilization_percent="75"] BOMThe Node IP "10.2.3.4" mount point "/" has reached a storage utilization of 75% with 2.371 GB free.`
Health Response	Response `{ Resource: /App/System/Storage?node_ip=10.2.3.4&mount_point=/ HQI { Color: Yellow Value: 1 } StatusText: The Node IP "10.2.3.4" mount point "/" has reached a storage utilization of 75% with 2.371 GB free. }`

Storage Utilization Full

31041	Storage Utilization Full
Description	Send an alert for each TPVM mount point when available storage is less than or equal to 1000 MB.
Preconditions	You cannot configure the default settings in the System Component (Monitor and System Service): • Polling frequency is **hourly** for the storage utilization threshold notice. • System monitors the TPVM storage utilization on the following mount points: ◦ **"/" (/dev/vda2)** ◦ **"/apps" (/dev/vdb1)** The polling service sends an alert with the TPVM storage utilization percentage.
Requirements	Alert shows the following data: • Node IP • Mount Point • Used MB • Available MB • Utilization Percent The following example shows an alert when a node IP mount point storage is 1000 MB or less (full): `<113>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"] [alert@1916 resource="/App/System/Storage?node_ip=10.2.3.4&mount_point=/" alertId="31041" cause="storageCapacityProblem" type="processingErrorAlarm" severity="critical"] [alertData@1916 node_ip="10.2.3.4" mount_point="/" used_mb="9485" available_mb="900" utilization_percent="98"] BOMThe Node IP "10.2.3.4" mount point "/" storage is full.`
Health Response	Response `{ Resource: /App/System/Storage?node_ip=10.2.3.4.&mount_point=/ HQI { Color: Black Value: 4 } StatusText: The Node IP "10.2.3.4" mount point "/" storage is full. }`

Storage Utilization Check

31042	Storage Utilization Check
Description	Send an alert for each TPVM mount point when capacity reaches below 75% of utilization.
Preconditions	You cannot configure default settings in the System Component (Monitor Service). • Polling frequency is **hourly** for storage utilization threshold notice. The system sends an info level storage utilization check alert when the monitor service starts up and the storage utilization is at a safe level and under the warning threshold. Alert severities that are higher than info level, are continually sent at the polling frequency.

Requirements	Alert shows the following data:
	• Node IP
	• Mount Point
	• Used MB
	• Available MB
	• Utilization Percent
	The following example shows an alert when a node IP mount point reaches below 75% of storage utilization:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]`
	`[alert@1916`
	`resource=”/App/System/Storage?node_ip=10.2.3.4&mount_point=/”`
	`alertId=”31042”`
	`cause=”storageCapacityCheck”`
	`type=”processingErrorAlarm”`
	`severity=”info”]`
	`[alertData@1916`
	`node_ip=”10.2.3.4”`
	`mount_point=”/”`
	`used_mb=”5114”`
	`available_mb=”4371”`
	`utilization_percent=”55”]`
	`BOMThe Node IP "10.2.3.4" mount point "/" is at a safe storage utilization of 55% with 4.371 GB free.`

	Response
	`{`
	`Resource: /App/System/Storage?node_ip=10.2.3.4&mount_point=/`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The Node IP "10.2.3.4" mount point "/" is at a safe storage utilization of 55% with 4.371 GB free.`
	`}`

Upgrade Alerts

Use the information in the following tables to learn about all possible upgrade alerts in detail that are raised by Fault Management.

XCO Upgrade Initiated

31060	XCO Upgrade Initiated
Description	Send an alert when XCO upgrade is initiated.
Preconditions	None

Requirements	Alert shows the following data:
	• Deployment Suite
	• Deployment Type
	• Deployment Platform
	• Original Version
	• New Version
	The following example shows an alert when an XCO upgrade is initiated:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]`
	`[alert@1916`
	`resource=””`
	`alertId=”31060”`
	`cause=”operationInitiated”`
	`type=”communicationsAlarm”`
	`severity=”info”]`
	`[alertData@1916`
	`deployment_suite=”fabric”`
	`deployment_type=”multi-node”`
	`deployment_platform=”TPVM”`
	`original_version=”3.2.0-v1”`
	`new_version=”3.2.0-v2”]`
	`BOMDeployment Suite "fabric" upgrade initiated for type "multi-node" on Platform "TPVM" with Original Version "3.2.0-v1" to New Version "3.2.0-v2".`

Health Response	Health resources are not associated with upgrade.

XCO Upgrade is successful

31061	XCO Upgrade is successful
Description	Send an alert when XCO upgrade is successful.
Preconditions	None

31061	XCO Upgrade is successful

Requirements	Alert shows the following data:
	• Deployment Suite
	• Deployment Type
	• Deployment Platform
	• Original Version
	• New Version
	
	The following example shows an alert when an XCO upgrade is successful:
	`Syslog RFC-5424 Example:`
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.20.30.40” enterpriseId=”1916” software=”XCO” swVersion=”3.2.0”]`
	`[alert@1916`
	`resource=””`
	`alertId=”31061”`
	`cause=”operationSuccess”`
	`type=”communicationsAlarm”`
	`severity=”info”]`
	`[alertData@1916`
	`deployment_suite=”fabric”`
	`deployment_type=”multi-node”`
	`deployment_platform=”TPVM”`
	`original_version=”3.2.0-v1”`
	`new_version=”3.2.0-v2”]`
	`BOMDeployment Suite "fabric" upgrade successfully completed for type "multi-node" on Platform "TPVM" with Original Version "3.2.0-v1" to New Version "3.2.0-v2".`

Health Response	Health resources are not associated with upgrade.

Missed Alerts

Use this topic to learn about the alerts when a service is down.

Fault Management Service Restart

The fault manager service maintains the alert sequence IDs and guarantees ordered sequencing of alert notifications even after the service reboots.

The incoming alerts remain on the message bus until the fault management service acknowledges it to remove it off the message bus.

Notification Service Restart

Fault Manager publishes alert notifications on the message bus. The notification service acknowledges all the alert notifications. If the notification service crashes or reboots, the un-acknowledged notifications that are left on the message bus, are published to the registered subscribers after the notification service reboots.

RabbitMQ Restart

XCO does not persist messages across MQ reboot, so all pending alerts are lost. You can query the fault service for the missed alert notifications using the sequence IDs.

Note

Depending on the state of the message location, fault management service might not receive notifications from the components. Therefore, it does not raise an alert.

There are minimal chance of RabbitMQ rebooting alone . RabbitMQ reboots are usually associated with system issues which impact other services.

System Restart

XCO attempts to re-notify fault management service on reboot. You can observe more frequent updates on HA status, storage status, and LDAP connectivity.

Most in-flight messages are lost. However, XCO ensures that alerts are regenerated and published on system restart. This also applies to failovers.

XCO increments the sequence ID correctly for alerts even after the system reboot and ensures ordered delivery of notifications.

Sub-System Restart

A sub-system is a component of XCO that is responsible for publishing a message that is eventually converted into an alert by the fault management service.

Note

XCO cannot raise alerts during a sub-system reboot.

Alert Order

The monitoring component service responsible for raising the HA failover alerts will publish the alert to fault manager once the necessary services are up and running. This delay in publishing might result in the HA failover alerts having out of order sequence numbers compared to other alerts.

The timestamp of the alert reflects when the alert has been raised by a component service.

Fault Management - Alarms

An alarm is a stateful entity that is raised and cleared solely by the system.

The alarm instance retains a list of time-stamped status changes (such as raising an alarm, updating the alarm severity, and clearing an alarm) as well as a list of time-stamped user state changes (such as acknowledging and closing an alarm). Both change lists have a limited circular buffer of 32 elements and the change history is maintained even after the closing and reopening of an alarm.

In XCO, a cleared alarm is considered an open alarm. You can acknowledge and close alarms or purge all closed alarms.

You can follow the typical alarm operations when investigating and administering alarms.

1. Acknowledge: You can acknowledge an open alarm with an optional comment to indicate that the alarm is being investigated. Additional user acknowledgments are allowed and tracked.
2. Close: You must close the alarm to remove it from the open alarms. You can provide an optional comment when closing an alarm. The system can raise or clear the alarm but the alarm is still considered opened until you close it. Once the alarm is closed, you cannot close it again or acknowledge it until the system opens and raises the alarm again.
3. Purge: All the closed alarms are removed and the associated resources are released.

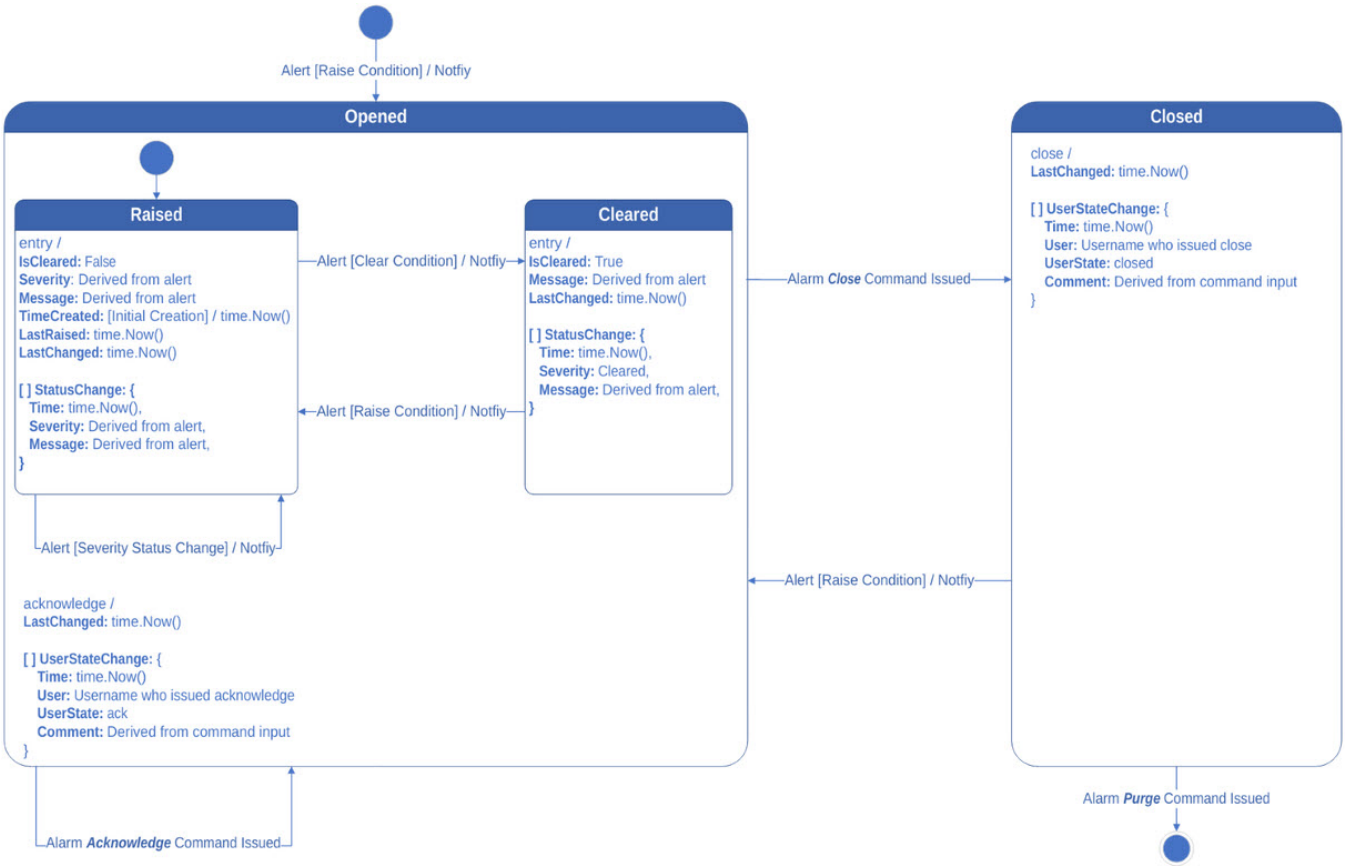

Figure 55: XCO alarm state transitions

Note

XCO supports a maximum of 500 alarms on TPVM deployments. You must purge the closed alarms manually to allow new alarms.

Alarm Inventory

The alarm inventory contains a list of system default alarms.

The following tables describe various parameters in alarm inventory.

Table 24: Alarm severity

Severity	Enum	Description
Critical	6	Service-affecting condition which requires an immediate corrective action.
Major	5	Service-affecting condition which requires an urgent corrective action.
Minor	4	Non-service-affecting condition which requires a corrective action.
Warning	3	Potential service-affecting fault which requires further diagnosis to prevent serious consequences.
Indeterminate	2	Severity level cannot be determined. **Note:** This severity must be avoided.
Cleared	1	Alarm is cleared by the system. **Note:** This severity level is only applicable to the “StatusChange” list severity.

Table 25: Alarm types

Types	Enum	Description
Other	1	Indicates a catch-all category for alarms. **Note:** Avoid this alarm type if other alarm types are applicable.
Communications	2	Indicates a communication related issue
QualityOfService	3	Indicates a quality related issue
Processing	4	Indicates a software processing issue
Equipment	5	Indicates faulty equipment
Environmental	6	Indicates an issue related to the enclosure housing the equipment
Integrity	7	Indicates duplicate, missing, modified, unexpected, or out of sequence information
Operational	8	Indicates that the provisioning of the requested service was not possible due to unavailability or malfunction of the service
Physical	9	Indicates cable tampering or intrusion

Table 25: Alarm types (continued)

Types	Enum	Description
Security	10	Indicates a security violation such as authentication failure or unauthorized access attempt
TimeDomain	11	Indicates that an event has occurred at an unexpected or prohibited time
Fabric	12	Indicates a fabric-service related issue
Tenant	13	Indicates a tenant-service related issue
Policy	14	Indicates a policy-service related issue

Table 26: Alarm inventory

Alarm Name	ID	Type	Max Alarm Instances	Description	Severity	Will Clear	Action
Certificate Expiration	32000	Security	7 (application certificate types)	Raise an alarm to notify that an XCO certificate is about to expire or has expired.	Warning / Critical	True	Notify APP_ALARMS syslog / webhook
Device Certificate Expiration	32001	Security	Number of Devices * 3 (device certificate types)	Raise an alarm to notify that a device certificate is about to expire or has expired.	Warning / Critical	True	Notify APP_ALARMS syslog / webhook

Table 26: Alarm inventory (continued)

Alarm Name	ID	Type	Max Alarm Instances	Description	Severity	Will Clear	Action
Login Authentication	32010	Security	Number of Users** Includes existing and non-existing application users.	Raise an alarm to notify suspicious login activity.	Warning (Only)	False	Notify APP_ALARMS syslog / webhook
Storage Utilization	32040	Processing	Number of Nodes * 2 (monitored mount points)	Raise an alarm to notify storage utilization for XCO installation has reached a certain threshold or it is full.	Warning - Critical (All severities between Warning and Critical are possible)	True	Notify APP_ALARMS syslog / webhook

Table 27: Associated alerts for alarms

Associated Alerts	Associated Alarms	Severity	Description
`CertificateExpiryNoticeAlert`	Certificate Expiration	Warning	Expiration of an XCO certificate within 90 days
`CertificateExpiredAlert`	Certificate Expiration	Critical	XCO certificate has expired
`CertificateRenewalAlert`	Certificate Expiration	Cleared	XCO certificate renewal
`DeviceCertificateExpiryNoticeAlert`	Device Certificate Expiration	Warning	Expiration of a device certificate within 90 days
`DeviceCertificateExpiredAlert`	Device Certificate Expiration	Critical	Device certificate has expired
`DeviceCertificateRenewalAlert`	Device Certificate Expiration	Cleared	Device certificate renewal
`LoginFailureAlert`	Login Authentication	Warning	5 successive failed login attempts within a minute **Note:** System does not clear the alarm.

Table 27: Associated alerts for alarms (continued)

Associated Alerts	Associated Alarms	Severity	Description
StorageUtilizationThresholdAlert	Storage Utilization	Warning	File system utilization is 75% - 84%
	Storage Utilization	Minor	File system utilization is 85% - 94%
	Storage Utilization	Major	File system utilization is 95% - 96%
	Storage Utilization	Critical	File system utilization is 97% - 100%
StorageUtilizationFullAlert	Storage Utilization	Critical	File system utilization is 97% - 100%
StorageUtilizationCheckAlert	Storage Utilization	Cleared	File system utilization is below 75%

Maximum Alarm Instance

The maximum alarm instance calculation is as follows:

- Number of devices: 20
- Number of users: 10 (existing and non-existing)
- Number of nodes: 2

```
CertificateExpiration Instances) + (DeviceCertificateExpiration Instances) +
(LoginAuthentication  Instances) + (StorageUtilization Instances) =
(7) + (20 * 3) + (10) + (2 * 2) = 81
```

Alarm Status Change Notifications

Alarms are responsible for sending notifications to any syslog and/or webhook subscribers subscribed to the APP_ALARMS notifications.

Alarm notifications are sent out when alarms are raised, cleared, and severities are updated.

Note
XCO 3.2.0 sends out similar notifications for APP_ALERTS and APP_ALARMS.

Table 28: Syslog Severity

Alarm Severity	Alert Severity	Syslog Severity	Syslog Enum	Description
		0	Emergency	System unusable
Critical	Critical	1	Alert	Immediate action required
Major	Major	2	Critical	Critical condition

Table 28: Syslog Severity (continued)

Alarm Severity	Alert Severity	Syslog Severity	Syslog Enum	Description
Minor	Minor	3	Error	Error condition
Warning	Warning	4	Warning	Warning condition
Indeterminate/ Cleared		5	Notice	Normal, but significant condition
	Info	6	Informational	Informational messages
		7	Debug	Debug-level messages

Table 29: Syslog Alarm (RFC-5674) - Common Alarm Payload

Field	SD-ID (Structured Data ID)	Example	Description
<###>	N/A	164 = (20 * 8) + 4 Alarm Range: 160-167	**Priority Value:** (Syslog Facility * 8) + Syslog Severity **Syslog Facility:** 20 local use 4 (XCO Alarms) See Table 28 on page 553.
Version	N/A	1	Version of syslog message
Timestamp	N/A	2003-10-11T22:14:15.003Z	Timestamp of syslog message
Hostname	N/A	xco.machine.com	Hostname of XCO
App Name	N/A	FaultManager	Application generating syslog alarm
Proc ID	N/A	-	Process ID
Msg ID	N/A	32000	Alarm sub-type classification
Sequence ID	meta	12	Tracks the sequence in which messages are submitted to the syslog transport. The APPS_ALARMS topic maintains its own sequence id compared to other topics.
IP	origin	10.20.30.40	IP address of XCO host
Enterprise ID	origin	1916	Extreme Networks Enterprise ID

Table 29: Syslog Alarm (RFC-5674) - Common Alarm Payload (continued)

Field	SD-ID (Structured Data ID)	Example	Description
Software	origin	XCO	Software Name
SW Version	origin	3.2.0	Software Version
Resource	alarm	/App/System/Security/Certificate?type=app_server_cert	XCO Health Resource path (with any query parameters) associated with the alarm.
ProbableCause	alarm	keyExpired	Reason for the Alarm (Attempt to map to IANA standards)
PerceivedSeverity	alarm	warning	Severity of the XCO Alarms See Table 24 on page 550.
EventType	alarm	security	Indicates the Category (Attempt to map to IANA standards)
BOMText	N/A	The application server certificate on the application will expire soon on “Sep 12 10:00:45 2022 GMT”.	(Byte Order Mask) Textual description of the Alarm’s status update.

The following is an example of Syslog Alarm:

```
<164>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - 32000

[meta sequenceId="12"]
[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.2.0"]
[alarm resource="/App/System/Security/Certificate?type=app_server_cert"
probableCause="keyExpired"
eventType="security"
perceivedSeverity="warning"]
[alarmData@1916
type="app_server_cert"
expiry_date="Sep 12 10:00:45 2022 GMT"]
BOMThe application server certificate on the application will expire soon on "Sep 12
10:00:45 2022 GMT".
```

The following is an example of Webhook Alarm:

```
{
  "type": "Alarm",
  "timestamp": "2003-10-11T22:14:15.003Z",
  "severity": "warning",
  "message": "The application server certificate on the application will expire soon on
\"Sep 12 10:00:45 2022 GMT\"",
  "application": "faultmanager",
  "source_ip": "10.20.30.40",
  "device_ip": "",
  "username": "",
  "message_id": "",
  "hostname": "tpvm1",
```

```
  "logtype": "",
  "task": "",
  "scope": "",
  "status": "",
  "sequence_id": 12,
  "alert_id": 0,
  "alarm_id": 32000,
  "resource": "/App/System/Security/Certificate?type=app_server_cert",
  "alarm_type": "security",
  "alarm_cause": "keyExpired",
  "alert_data": null,
  "alarm_data": {
    "type": "app_server_cert",
    "expiry_date": "Sep 12 10:00:45 2022 GMT",
  }
}
```


Alarm Commands

XCO 3.2.0 supports the following system alarm commands:

Commands	Description
`efa system alarm inventory`	Lists supported alarms
`efa system alarm show`	Lists active alarms
`efa system alarm summary`	Lists alarm counts and statistics
`efa system alarm acknowledge`	Acknowledges the alarm instance
`efa system alarm close`	Closes an open alarm
`efa system alarm purge`	Purges all closed alarms

For more information, see *ExtremeCloud Orchestrator Command Reference, 3.2.1*.

Health Management

Health score provides a quick mechanism to figure out the overall health of XCO and its components. The entire health score is represented as a hierarchy which can be accessed using the Resource Path.

You can query health at various levels by providing the appropriate Resource Path.

Note
Alarms do not directly affect system health.

Bubbling of Health Status

In the following example, the certificate present in `/App/System/Security/Certificate` directory is in yellow state and goes up to the top level of the hierarchy.

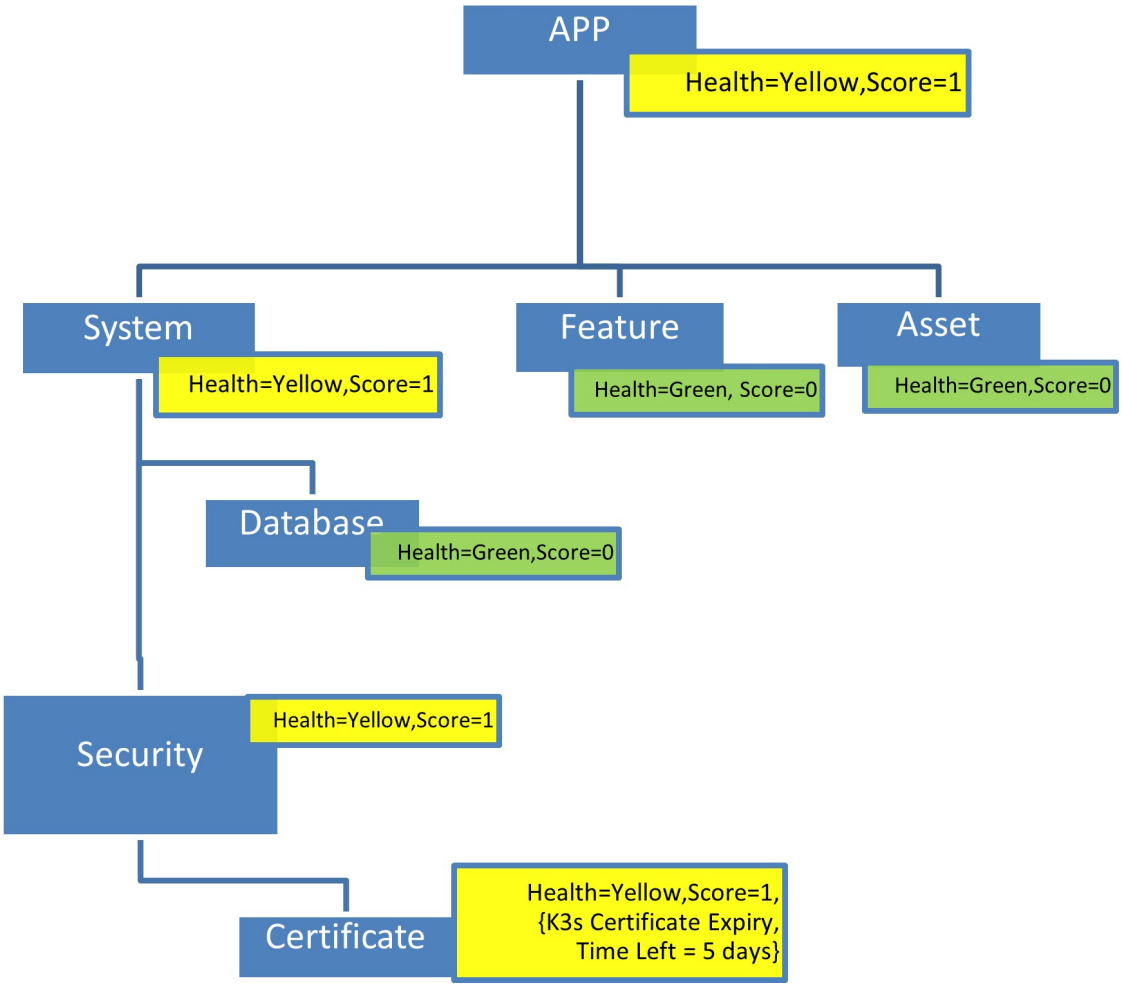
APP
Health=Yellow,Score=1
System
Health=Yellow,Score=1
Feature
Health=Green, Score=0
Asset
Health=Green,Score=0
Database
Health=Green,Score=0
Security
Health=Yellow,Score=1
Certificate
Health=Yellow,Score=1,
{K3s Certificate Expiry,
Time Left = 5 days}

Each node in the hierarchy maintains the following details:

Node Details	Description
Health Status and Score	It represents the Health Quality indicator (HQI) of the system The Health Quality indicator has Color and Value which is described in the next table. The HQI indicator goes up in the hierarchical system.
Additional Metadata	Every node can optionally maintain metadata about what provides the reason for the current HQI.

HQI Color	HQI Value	Description
Green	0	Healthy System
Yellow	1	Potential Failure is imminent For example, certificate expiry in 10 days.
Orange	2	Failed (Yet System Functional) For example, standby node is down.
Red	3	Requires Immediate Attention For example, database server is down.
Black	4	Critical

Health Commands

XCO supports the following health commands:

Commands	Description
`efa health inventory show`	List of health resources
`efa health show`	List of top level health resources
`efa health show –resource [ App/System/Security/Certificate ]`	Health score for a resource
`efa health detail show –resource [ /App/System/Security/Certificate ]`	Health score and details at the specified resource

Health APIs

API	Description	
`Get Health [<Resource>]`	Gets the HQI (Health QualityIndicator) of XCO	
	`Get Health`	Gets Overall HQI Values
	`Get Health /App/System`	Gets HQI values at System Level
`Get Health Detail <Resource>`	`Gets Health`	
	`Get Health Detail /App/System/Certificate`	Gets HQI values and metadata at certificate

Example

```
Request Get Health

Response
{
    Resource: /App
    HQI {
        Color: Yellow
        Value: 1
    }
    StatusText: [<Freeform status text>]
}
Request Get Health Detail

Response
{
    Resource: /App
    HQI {
        Color: Yellow
        Value: 1
    }
   Contributor {
        ResourceList: [/App/System/Certificate, /App/Component/Fabric]
   }
   StatusText: [<Freeform status text>]
}
Request Get Health /App/System/Certificate

Response
{
    Resource:/App/System/Certificate
    HQI {
        Color: Yellow
        Value: 1
    }
    StatusText: [Certificate x expires on <date>.]
}
```

Health service URI has additional query parameters to enable users to query for specific type of entity and particular asset. As part of the API validation the query parameters are validated and appropriate response is returned. Check the REST API guide for the different status and responses.

Fabric Health

Fabric health shows the current health state of a fabric.

The state of a fabric health can be either Black (Critical), Red (Degraded) or Green (Healthy). It is derived based on the fabric status, fabric level physical topology health, and device health.

The following factors decides the fabric health:

- **Critical**: It takes the highest priority. If any status is critical, the fabric health will also be critical.
- **Degraded**: It takes the next priority. If any status is degraded, the fabric health will also be degraded.

- **Healthy**: If all the statuses are in healthy state, the fabric health will be healthy.

Note

- Green status indicates a healthy system.
- Red status indicates that the system is degraded and an immediate attention is required.
- Black status indicates that the system condition is critical.
- If no devices are added to fabric, then the fabric health is Green.

Fabric Health Calculation

The following diagram shows the fabric health in a hierarchical format. Any change in a child health results in the propagation of health to the parent in the hierarchy.

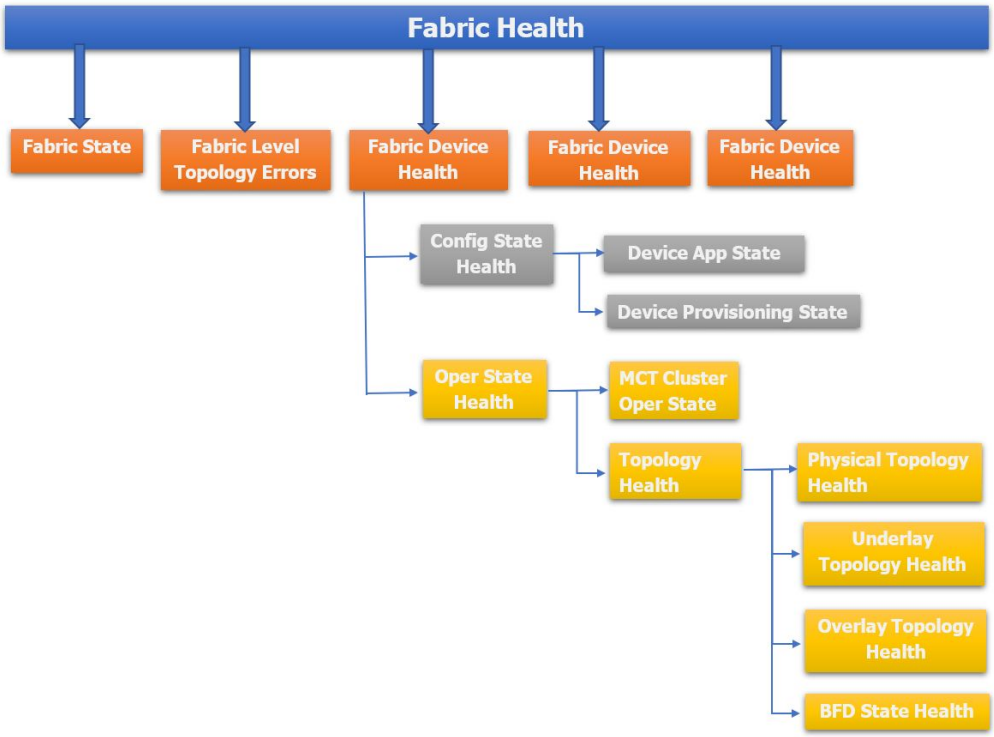

The following is the truth table for a fabric health:

Fabric Status	Fabric Level Physical Topology Health	Device Health	Fabric Health
Configure-success	Green	Green	Green
Configure-success	Green	Red	Red
Configure-success	Green	Black	Black
Configure-success	Red	Green	Red
Configure-success	Red	Red	Red

Fabric Status	Fabric Level Physical Topology Health	Device Health	Fabric Health
Configure-success	Red	Black	Black
Configure-success	Black	Green, Red, Black	Black
Configure-success	Black		
Configure-success	Black		
created	Green	Green	Red
created	Green	Red	Red
created	Green	Black	Black
Created	Red	Green	Red
Created	Red	Red	Red
Created	Red	Black	Black
Created	Black	Green, Red, Black	Black
Migrate-success	Green	Green	Red
Migrate-success	Green	Red	Red
Migrate-success	Green	Black	Black
Migrate-success	Red	Green	Red
Migrate-success	Red	Red	Red
Migrate-success	Red	Black	Black
Migrate-success	Black	Green, Red, Black	Black
Migrate-Failure	Green	Green	Red
Migrate-Failure	Green	Red	Red
Migrate-Failure	Green	Black	Black
Migrate-Failure	Red	Green	Red
Migrate-Failure	Red	Red	Red
Migrate-Failure	Red	Black	Black
Migrate-Failure	Black	Green, Red, Black	Black
Setting-updated	Green	Green	Red
Setting-updated	Green	Red	Red
Setting-updated	Green	Black	Black
Setting-updated	Red	Green	Red
Setting-updated	Red	Red	Red
Setting-updated	Red	Black	Black
Setting-updated	Black	Green, Red, Black	Black
created	Green	Green	Red
created	Green	Red	Red
configure-failed	Green, Black, Red	Green, Black, Red	Black

Fabric Status

The following table shows the mapping of a fabric status to the fabric health.

Fabric Status	Health
created	Green
configure-success	Green
configure-failed	Black
settings-updated	Red
migrate-success	Red
migrate-failed	Red

Fabric Level Physical Topology Health

The fabric level physical topology health is calculated based on the fabric level errors of the topology.

The Physical topology health is divided into two parts:

- Fabric level physical topology health: It is used to calculate a fabric health.
- Device level physical topology health: It is used to calculate a device operational state health.

Fabric Level Physical Topology Validations for non-Clos Fabric

- Ensure that each rack must have two devices.

Non-Clos Fabric Level Topology Error Scenarios

- Only one device is present in the rack

Fabric Level Physical Topology Validations for Clos Fabric

1. Ensure that the stage 3 fabric contains at least one leaf or border leaf and spine device.
2. Ensure that the stage 5 fabric contains at least one leaf or border leaf and super-spine devices

Clos Fabric Level Topology Error Scenarios

- For 3-stage Clos fabric:
  1. No leaf or border leaf devices in fabric
  2. No spine devices in the fabric
- For 5-stage Clos fabric:
  1. No leaf or border leaf devices in fabric
  2. No Super-spine devices in the fabric

If the topology is invalid, the fabric level physical topology health gets degraded or else it is healthy.

Device Health

You can calculate a device health based on the configuration and operational status of device.

The following table describes the device health based on its configuration and operational state health:

Device	Config state health	Operational state health	Device health
D1	Green	Green	Green
D2	Green	Red	Red
D3	Red	Red	Red
D4	Red	Black	Black
D5	Black	Red	Black
D6	Black	Black	Black

Config State Health

You can calculate a config state health based on the app and device state.

The following table provides application states and the corresponding configuration states health:

App State	Health
cfg ready	Red
cfg in-sync	Green
cfg error	Black
cfg refreshed	Red
cfg refresh error	Black
cfg unknown	Red
device remove failed	Black

The following table provides health of the device state based on its provisioning status:

Dev State	Health
not provisioned	Red
provisioned	Green
provisioning failed	Black
unknown	Red

Operational State Health

You can determine an operational health status of a device based on the cluster health, device level physical topology health, device level underlay topology health, and the device level overlay topology health.

The following table describes the operational health of devices based on the physical and underlay topology and cluster health:

Device IP	Cluster Health	Device Physical Topology Health	Device Underlay Topology Health	Operational State Health
D1	Green	Green	Green	Green
D2	Green	Red	Red	Red
D3	Black	Green	Green	Black
D4	Black	Red	Red	Black
D5	Green	Green	Black	Black
D6	Black	Black	Black	Black

Cluster Health

The cluster health is obtained from the peer, cluster, and the peer keep-alive state. You can use the **`show cluster`** command output of a switching device to view these states.

The following table describes the health of a cluster:

Peer State	Cluster State	Peer Keep-alive State	Cluster Health
Up	Up	Up	Green
Up	Up	Down	Red
Up	Down	Up	Red
Up	Down	Down	Red
Down	Up	Up	Black
Down	Up	Down	Black
Down	Down	Up	Black
Down	Down	Down	Black

Device Level Physical Topology Health

Use this topic to learn about health of a device level physical topology.

If a topology is invalid, the device level physical topology health is degraded or else healthy.

Device level physical topology validations for non-Clos fabric

- Ensure that two devices in rack must have link between them.
- Ensure that each rack is connected to at least another rack.

Device level physical topology validations for Clos fabric

- Leaf node must be connected to all the Spine nodes.
- Spine node must be connected to all the Leaf nodes.
- Border Leaf node must be connected to all the Spine nodes or Superspine nodes but not both.
- Spine node must be connected to all the Border Leaf nodes.
- No more than two Leaf nodes must not be connected to each other.
- No more than two Border Leaf nodes must not be connected to each other.
- Border leaf node and leaf node must not be connected.
- Spine nodes must not be connected to each other.
- Spine nodes must not be connected to each other.
- Super Spine nodes must not be connected to each other
- If a Leaf node is marked as "multi-homed", then the node must have a MCT neighbor.
- If a Leaf node is marked as "single-homed", then the node must not be connected to other Leaf node(s).
- If a Border Leaf node is marked as "multi-homed", then the node must have a MCT neighbor.
- If a Border Leaf node is marked as "single-homed", then the node must not be connected to other Border Leaf node(s).

Device Level Underlay Topology Health

Learn about health of a device level underlay topology.

- An ESTABILISHED underlay session state is considered as a Green (Healthy) state.
- When the underlay session state of all the BGP neighbors is in ESTABLISHED state, the underlay topology health is Green (Healthy).
- When a BGP neighbor (between the two devices) is not in ESTABLISHED state, the underlay topology health is Red (Degraded).
- When all the neighbors between the devices are down, the underlay topology health is Black (Critical).

The following table describes device level underlay health based on BGP neighbors and underlay session state:

Device	Remote Device	BGP Neighbor Configured on Device Pointing to Remote Device	Underlay Session State	Device Underlay Health
Device1	Device2	Neighbor1	ESTAB	Green
		Neighbor2	ESTAB	
	Device3	Neighbor3	ESTAB	
		Neighbor4	ESTAB	

Device	Remote Device	BGP Neighbor Configured on Device Pointing to Remote Device	Underlay Session State	Device Underlay Health
Device 4	Device5	Neighbor5	CONN	Red
		Neighbor6	ESTAB	
	Device6	Neighbor7	ESTAB	
		Neighbor8	ESTAB	
Device7	Device8	Neighbor10	CONN	Black
		Neighbor11	CONN	
	Device9	Neighbor12	ESTAB	

Device Level Overlay Topology Health

Learn about health of a device level overlay topology. An overlay topology is defined after the tunnels are created in a fabric.

- Overlay topology health is applicable only for leaf or border leaf devices.
- When all the tunnels are up, then the overlay topology health is healthy. If one of the tunnels is not up, then it is critical (black).

The following table describes device level overlay health based on admin and operational state of tunnels:

		Tunnels	Admin State	Oper State	Device Overlay Health
Device1	Device2	Tunnel1	up	up	Green
			up	up	
	Device3	Tunnel2	up	up	
			up	up	
Device 4	Device5	Tunnel3	up	down	Black
			up	down	
	Device6	Tunnel4	up	up	
			up	up	
Device7	Device8	Tunnel5 (not configured)			Black
	Device9	Tunnel6	up	up	

Sample Output of 3-stage Clos Fabric Creation

The following sections show sample output from the operations involved in creating a 3-stage Clos fabric. It normally takes between 1 and 15 minutes to generate this output for a device.

Create 3-Stage Clos Fabric

The following sample output creates a 3-stage Clos fabric:

```
(efa:user)user@dev-server:~$ efa fabric create --name fab3
Create Fabric fab3 [Success]

--- Time Elapsed: 67.967803ms ---
(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: Green, Fabric Health: Red
+-------+---+----+---+----+------+-----+----------+-------+--------+
|IP     |POD|HOST|ASN|ROLE|DEVICE|APP  |CONFIG    |PENDING|VTLB| LB|
|ADDRESS|   |NAME|   |    |STATE |STATE|GEN REASON|CONFIG |ID  | ID|
+-------+---+----+---+----+------+-----+----------+-------+----+---+
+-------+---+----+---+----+------+-----+----------+-------+----+---+

REST API Get Fabric:
http://10.37.35.163/v1/fabric/fabric?name=fab3
{
    "fabric-name": "fab3",
    "fabric-id": 7,
    "fabric-stage": 3,
    "fabric-type": "clos",
    "fabric-status": "Green",
    "fabric-health": "Red",
    "fabric-settings": {
        "AllowASIn": "0",
        "AnyCastMac": "0201.0101.0101",
        "BFDEnable": "Yes",
        "BFDMultiplier": "3",
        "BFDRx": "300",
        "BFDTx": "300",
        "BackupRoutingEnable": "No",
        "BackupRoutingIpv4Range": "10.40.40.0/24",
        "BackupRoutingIpv6Range": "fd40:4040:4040:1::/120",
        "BgpDynamicPeerListenLimit": "100",
        "BorderLeafASNBlock": "66000-66100",
        "ConfigureOverlayGateway": "Yes",
        "ControlVE": "4090",
        "ControlVlan": "4090",
        "DefaultMdtgroup": "239.1.1.1",
        "DuplicateMacTimer": "5",
        "DuplicateMaxTimerMaxCount": "3",
        "IPMTU": "9100",
        "IPV6AnyCastMac": "0201.0101.0102",
        "LacpTimeout": "long",
        "LeafASNBlock": "65000-65534",
        "LeafPeerGroup": "spine-group",
        "LoopBackIPRange": "172.31.254.0/24",
        "LoopBackPortNumber": "1",
        "MCTLinkIPRange": "10.20.20.0/24",
        "MTU": "9216",
        "MacAgingConversationalTimeOut": "300",
        "MacAgingConversationalTimeout": "300",
        "MacAgingTimeout": "1800",
        "MacMoveLimit": "20",
        "MaxPaths": "8",
        "MctPortChannel": "64",
        "Md5PasswordEnable": "No",
        "MdtgroupRange": "239.0.0.0/8",
        "OptimizedReplicationEnable": "No",
```

```
        "OverlayGwBroadcastLocalBiasEnable": "No",
        "P2PIPType": "numbered",
        "P2PLinkRange": "10.10.10.0/23",
        "SpineASNBlock": "64512-64768",
        "SpinePeerGroup": "leaf-group",
        "SuperSpineASNBlock": "64769",
        "SuperSpinePeerGroup": "spine-group",
        "VNIAutoMap": "Yes",
        "VTEPLoopBackPortNumber": "2"
    },
    "number-of-pods": "0",
    "number-of-racks": "0",
    "number-of-single-homed-leaf-nodes": "0",
    "number-of-multi-homed-leaf-nodes": "0",
    "number-of-spine-nodes": "0",
    "number-of-single-homed-border-leaf-nodes": "0",
    "number-of-multi-homed-border-leaf-nodes": "0",
    "number-of-super-spine-nodes": "0",
    "number-of-not-provisioned-nodes": "0",
    "number-of-provisioned-nodes": "0",
    "number-of-provisioned-failed-nodes": "0",
    "number-of-config-ready-nodes": "0",
    "number-of-config-generation-error-nodes": "0",
    "number-of-config-in-sync-nodes": "0",
    "number-of-config-refreshed-nodes": "0",
    "fabric-devices": {}
}
```

Add Two Spine Devices 10,20.246.1 and 2

The following sample output configures two spine devices in a 3-stage Clos fabric:

```
(efa:user)user@dev-server:~$ efa fabric device add --role spine --ip 10.20.246.2 --name
fab3
Inventory Device(s) Registration[success]
+----+-------------+-----------+-------+--------------+------------------------------+
| ID | IP Address  | Host Name | Model | Chassis Name |           Firmware           |
+----+-------------+-----------+-------+--------------+------------------------------+
| 7  | 10.20.246.2 | NH-2      | 3012  | SLX9250-32C  | 20.4.3slxos20.4.3_221117_0600 |
+----+-------------+-----------+-------+--------------+------------------------------+
Device Details
Add Device(s) [Success]

        Addition of Spine device with ip-address = 10.20.246.2 [Succeeded]
Validate Fabric [Failed]
        No Leaf Devices
Error :  fabric validation failed

Add device to fabric



(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
==============
Fabric Name                           :  fab3
Fabric Type                           :  clos
Fabric Health                         :  Red
 Fabric Status                        :  Green
 Fabric Level Physical Topology Health  :  Red
Fabric Device Health
+-------------+-------+--------------------+------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+--------------------+------------------+---------------+
```

```
| 10.20.246.1 | Spine | reg                 | Green             | Red           |
| 10.20.246.2 | Spine | Red                 | Green             | Red           |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
================
--- Time Elapsed: 32.248825ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Red
 Fabric Status                           :  Green
 Fabric Level Physical Topology Health   :  Red
------------------------------------------------------------------------------------------------
----------------
Fabric level topology errors     :
+---------------------+----------------+---------------+
| MISSING SUPERSPINES | MISSING SPINES | MISSING LEAFS |
+---------------------+----------------+---------------+
| true                | false          | true          |
+---------------------+----------------+---------------+
------------------------------------------------------------------------------------------------
----------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.1 [Spine]
  Device Health                          :  Red
  Configuration State Health             :  Red
    Dev State                            :  not provisioned
    App State                            :  cfg ready
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.2 [Spine]
  Device Health                          :  Red
  Configuration State Health             :  Red
    Dev State                            :  not provisioned
    App State                            :  cfg ready
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
------------------------------------------------------------------------------------------------
-----------------
```

Add Two MCT Pairs 10.20.246.5,6 and 10.20.20.246.3 and 4

The following sample output configures two MCT pairs in a 3-stage Clos fabric:

```
(efa:user)user@dev-server:~$ efa fabric device add-bulk --leaf
10.20.246.5,10.20.246.6,10.20.246.3,10.20.246.4 --name fab3
Inventory Device(s) Registration[Success]
+----+-------------+-----------+-------+--------------+-------------------------------
+--------+--------+
| ID | IP Address  | Host Name | Model | Chassis Name |
Firmware            | Status | Reason |
+----+-------------+-----------+-------+--------------+-------------------------------
+--------+--------+
| 11 | 10.20.246.3 | NH-Leaf1  | 3012  | SLX9250-32C  | 20.4.2slxos20.4.2_220803_1000
|        |        |
+----+-------------+-----------+-------+--------------+-------------------------------
+--------+--------+
```

```
| 9  | 10.20.246.4 | NH-Leaf2  | 3012  | SLX9250-32C  | 20.4.2slxos20.4.2_220803_1000
|        |        |
+----+-------------+-----------+-------+--------------+------------------------------
+--------+--------+
| 3  | 10.20.246.5 | NHF-Leaf1 | 3009  | SLX9150-48Y  | 20.4.1b
|        |        |
+----+-------------+-----------+-------+--------------+------------------------------
+--------+--------+
| 1  | 10.20.246.6 | NHF-Leaf2 | 3009  | SLX9150-48Y  | 20.4.1b
|        |        |
+----+-------------+-----------+-------+--------------+------------------------------
+--------+--------+
Device Details
Updating devices that are already registered: [10.20.246.3 10.20.246.4
10.20.246.5 10.20.246.6]
        Inventory Update with ip-address = 10.20.246.5 [Succeeded]
        Inventory Update with ip-address = 10.20.246.6 [Succeeded]
        Inventory Update with ip-address = 10.20.246.4 [Succeeded]
        Inventory Update with ip-address = 10.20.246.3 [Succeeded]
Add Device(s) [Success]

        Addition of Leaf device with ip-address = 10.20.246.6 [Succeeded]
        Addition of Spine device with ip-address = 10.20.246.2 [Succeeded]
        Addition of Leaf device with ip-address = 10.20.246.5 [Succeeded]
        Addition of Leaf device with ip-address = 10.20.246.3 [Succeeded]
        Addition of Spine device with ip-address = 10.20.246.1 [Succeeded]
        Addition of Leaf device with ip-address = 10.20.246.4 [Succeeded]
Validate Fabric [Success]




(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type:
clos, Fabric Status: Green, Fabric Health: Red
+-------------+-----+-----------+-------+-------+-----------------+-----------
+-------------------+-------------------------------------------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  |  DEVICE STATE   | APP STATE
| CONFIG GEN REASON |               PENDING CONFIGS                | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+-----------------+-----------
+-------------------+-------------------------------------------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | not provisioned | cfg ready
| DD,DA             | SYSP-C,BGP-C,INTIP-C                            | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | not provisioned | cfg ready
| DD,DA             | SYSP-C,BGP-C,INTIP-C                            | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65000 | leaf  | not provisioned | cfg ready
| DA                | SYSP-C,MCT-C,MCT-PA,BGP-C,INTIP-C,EVPN-C,O-C | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65000 | leaf  | not provisioned | cfg ready
| DA                | SYSP-C,MCT-C,MCT-PA,BGP-C,INTIP-C,EVPN-C,O-C | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | not provisioned | cfg ready
| DA                | SYSP-C,MCT-C,MCT-PA,BGP-C,INTIP-C,EVPN-C,O-C | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | not provisioned | cfg ready
| DA                | SYSP-C,MCT-C,MCT-PA,BGP-C,INTIP-C,EVPN-C,O-C | 2       | 1     |
+-------------+-----+-----------+-------+-------+-----------------+-----------
+-------------------+-------------------------------------------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface
Add/Delete/Update, PLC/PLD/PLU - IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway
Delete/Update, EU/ED - Evpn Delete/Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device
Add/ReAdd, ASN - Asn Update, SYS - System Properties Update
```

```
MD5 - BGP MD5 Password, BGPU - Router BGP
Update, BGPLL - BGP Listen Limit, POU - Port Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway,
SYSP - System Properties, INTIP - Interface IP, BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 155.196132ms ---
```

When you add spine and leaf devices, the fabric level physical topology state becomes Green. But the device health will be in Red state as cluster and BGP neighbors are not configured on the device. These are device level errors.

```
(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Red
 Fabric Status                           :  Green
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+----------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+----------------+
| 10.20.246.6 | Leaf  | Red                 | Red               | Red           |
| 10.20.246.5 | Leaf  | Red                 | Red               | Red           |
| 10.20.246.1 | Spine | Red                 | Red               | Red           |
| 10.20.246.2 | Spine | Red                 | Red               | Red           |
| 10.20.246.4 | Leaf  | Red                 | Red               | Red           |
| 10.20.246.3 | Leaf  | Red                 | Red               | Red           |
+-------------+-------+---------------------+-------------------+----------------+

================================================================================================
================================================================
--- Time Elapsed: 46.271008ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Red
 Fabric Status                            :  Green
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Red
  Configuration State Health              :  Red
    Dev State                             :  not provisioned
    App State                             :  cfg ready
  Operational State Health                :  Red
    Cluster Health                        :  Red
      Operational State                   :  false
```

```
      Peer Operational State            :  false
      Peer Keepalive Operational State  :  false
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Red
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.6      | 10.20.246.5           | 172.31.254.92           | 10.20.20.13 |
65001             | 65001                 | default-vrf | ipv4              |
unicast       |               | neighbor_not_configured |
| 10.20.246.6      | 10.20.246.2           | 172.31.254.92           | 10.10.10.27 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          |               | neighbor_not_configured |
| 10.20.246.6      | 10.20.246.2           | 172.31.254.92           | 10.10.10.27 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       |               | neighbor_not_configured |
| 10.20.246.6      | 10.20.246.1           | 172.31.254.92           | 10.10.10.25 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          |               | neighbor_not_configured |
| 10.20.246.6      | 10.20.246.1           | 172.31.254.92           | 10.10.10.25 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       |               | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                     :  10.20.246.5 [Leaf]
  Device Health                       :  Red
  Configuration State Health          :  Red
    Dev State                         :  not provisioned
    App State                         :  cfg ready
  Operational State Health            :  Red
    Cluster Health                    :  Red
      Operational State               :  false
      Peer Operational State          :  false
      Peer Keepalive Operational State  :  false
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Red
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.5      | 10.20.246.6           | 172.31.254.196          | 10.20.20.12 |
65001             | 65001                 | default-vrf | ipv4              |
unicast       |               | neighbor_not_configured |
| 10.20.246.5      | 10.20.246.1           | 172.31.254.196          | 10.10.10.29 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          |               | neighbor_not_configured |
| 10.20.246.5      | 10.20.246.1           | 172.31.254.196          | 10.10.10.29 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       |               | neighbor_not_configured |
```

```
| 10.20.246.5      | 10.20.246.2           | 172.31.254.196          | 10.10.10.31 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          |                | neighbor_not_configured |
| 10.20.246.5      | 10.20.246.2           | 172.31.254.196          | 10.10.10.31 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       |                | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ---------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Red
  Configuration State Health           :  Red
    Dev State                          :  not provisioned
    App State                          :  cfg ready
  Operational State Health             :  Red
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Red
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.1      | 10.20.246.4           | 172.31.254.144          | 10.10.10.16 |
64512             | 65000                 | default-vrf | l2vpn             |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.4           | 172.31.254.144          | 10.10.10.16 |
64512             | 65000                 | default-vrf | ipv4              |
unicast       |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.3           | 172.31.254.144          | 10.10.10.20 |
64512             | 65000                 | default-vrf | l2vpn             |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.3           | 172.31.254.144          | 10.10.10.20 |
64512             | 65000                 | default-vrf | ipv4              |
unicast       |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.6           | 172.31.254.144          | 10.10.10.24 |
64512             | 65001                 | default-vrf | l2vpn             |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.6           | 172.31.254.144          | 10.10.10.24 |
64512             | 65001                 | default-vrf | ipv4              |
unicast       |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.5           | 172.31.254.144          | 10.10.10.28 |
64512             | 65001                 | default-vrf | l2vpn             |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.5           | 172.31.254.144          | 10.10.10.28 |
64512             | 65001                 | default-vrf | ipv4              |
unicast       |                | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ---------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Red
  Configuration State Health           :  Red
    Dev State                          :  not provisioned
    App State                          :  cfg ready
  Operational State Health             :  Red
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Red
```

```
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.2      | 10.20.246.4           | 172.31.254.205          | 10.10.10.18 |
64512             | 65000                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.4           | 172.31.254.205          | 10.10.10.18 |
64512             | 65000                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.3           | 172.31.254.205          | 10.10.10.22 |
64512             | 65000                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.3           | 172.31.254.205          | 10.10.10.22 |
64512             | 65000                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.6           | 172.31.254.205          | 10.10.10.26 |
64512             | 65001                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.6           | 172.31.254.205          | 10.10.10.26 |
64512             | 65001                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.5           | 172.31.254.205          | 10.10.10.30 |
64512             | 65001                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.2      | 10.20.246.5           | 172.31.254.205          | 10.10.10.30 |
64512             | 65001                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Red
  Configuration State Health           :  Red
    Dev State                          :  not provisioned
    App State                          :  cfg ready
  Operational State Health             :  Red
    Cluster Health                     :  Red
      Operational State                :  false
      Peer Operational State           :  false
      Peer Keepalive Operational State  :  false
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Red
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.4      | 10.20.246.3           | 172.31.254.90           | 10.20.20.10 |
65000             | 65000                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.4      | 10.20.246.1           | 172.31.254.90           | 10.10.10.17 |
```

```
65000             | 64512                  | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.4      | 10.20.246.1           | 172.31.254.90           | 10.10.10.17 |
65000             | 64512                  | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.4      | 10.20.246.2           | 172.31.254.90           | 10.10.10.19 |
65000             | 64512                  | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.4      | 10.20.246.2           | 172.31.254.90           | 10.10.10.19 |
65000             | 64512                  | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.3 [Leaf]
  Device Health                         :  Red
  Configuration State Health            :  Red
    Dev State                           :  not provisioned
    App State                           :  cfg ready
  Operational State Health              :  Red
    Cluster Health                      :  Red
      Operational State                 :  false
      Peer Operational State            :  false
      Peer Keepalive Operational State  :  false
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Red
    Device underlay topology errors
+------------+-------------+-----------------+-------------+-----------+-------------
+-------------+----------+---------+----------+-------------------------+
| SOURCE     | DESTINATION | SOURCE DEVICE   | NEIGHBOR IP | SOURCE    | DESTINATION
|     VRF     | NEIGHBOR | NEIGHBOR| UNDERLAY |          ERROR          |
| DEVICE IP  | DEVICE IP   | DEVICE ROUTER ID|             | DEVICE ASN| DEVICE ASN
|             | AFI STATE| SAFI    | STATE    |                         |
+------------+-------------+-----------------+-------------+-----------+-------------
+-------------+----------+---------+----------+-------------------------+
| 10.20.246.3| 10.20.246.4 | 172.31.254.145  | 10.20.20.11 | 65000     | 65000       |
default-vrf | ipv4     | unicast |          | neighbor_not_configured |
| 10.20.246.3| 10.20.246.1 | 172.31.254.145  | 10.10.10.21 | 65000     | 64512       |
default-vrf | l2vpn    | evpn    |          | neighbor_not_configured |
| 10.20.246.3| 10.20.246.1 | 172.31.254.145  | 10.10.10.21 | 65000     | 64512       |
default-vrf | ipv4     | unicast |          | neighbor_not_configured |
| 10.20.246.3| 10.20.246.2 | 172.31.254.145  | 10.10.10.23 | 65000     | 64512       |
default-vrf | l2vpn    | evpn    |          | neighbor_not_configured |
| 10.20.246.3| 10.20.246.2 | 172.31.254.145  | 10.10.10.23 | 65000     | 64512       |
default-vrf | ipv4     | unicast |          | neighbor_not_configured |
+------------+-------------+-----------------+-------------+-----------+-------------
+-------------+----------+---------+----------+-------------------------+
-----------------------------------------------------------------------------------------------
-------------------------------------------------------------------

===============================================================================================
=====================================================================
```

Configure the Fabric

The following sample output configures a 3-stage Clos fabric:

```
(efa:user)user@dev-server:~$ efa fabric configure --name fab3
Validate Fabric [Success]
Configure Fabric [Success]
Please verify the fabric physical/underlay topology using 'efa fabric topology show
{physical | underlay}' before attempting tenant configuration on the fabric.
--- Time Elapsed: 1m56.314230141s ---
```

```
After this command it will take around 1 minute to move to healthy as cluster health, bgp
session states must get updated from device.

(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: healthy
+-------------+-----+-----------+-------+-------+-------------+-------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+-------------+-------------
+-------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+-------------+-------------
+-------------------+-----------------+---------+-------+


Fabric health is healthy as all the child contributors are healthy


(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+----------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+----------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+----------------+

================================================================================================
=====================================================================
--- Time Elapsed: 43.144657ms ---

(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
```

```
 Fabric Level Physical Topology Health   :  Green
-------------------------------------------------------------------------------------------------
-----------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                       :  10.20.246.6 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.5 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.1 [Spine]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.2 [Spine]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.4 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.3 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
-------------------------------------------------------------------------------------------------
-----------------------------------------------------------------------
```

```
================================================================================================
================================================================================
--- Time Elapsed: 37.895545ms ---
```

Shut Down the MCT Port Channel on Leaf Devices

The following sample output shuts down an MCT port channel on leaf devices:

```
Shutdown port channel of 10.20.246.5(mct pair 10.20.246.6) and 10.20.246.3(mct pair
10.20.246.4)

NHF-Leaf1(config)# interface Port-channel 64
NHF-Leaf1(config-Port-channel-64)# shutdown
NHF-Leaf1(config-Port-channel-64)#

NH-Leaf1(config)# interface Port-channel 64
NH-Leaf1(config-Port-channel-64)# shutdown
NH-Leaf1(config-Port-channel-64)#

Need to wait for some time to get the raslog events and update in fabric



(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |     APP STATE     |
CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync       |
NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync       |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg refresh error |
LD,IU,POU         | BGP-U,INTIP-U   | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg refresh error |
LD,POU            |                 | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg refresh error |
LD,IU,POU         | BGP-U           | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg refresh error |
LD,POU            |                 | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+-----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete
```

```
For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 149.741541ms ---



Fabric Health has moved to Black state due to port channel shutdown and detailed output
is as below

(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+----------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+----------------+
| 10.20.246.6 | Leaf  | Black               | Black             | Black         |
| 10.20.246.5 | Leaf  | Black               | Black             | Black         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Black               | Black             | Black         |
| 10.20.246.3 | Leaf  | Black               | Black             | Black         |
+-------------+-------+---------------------+-------------------+----------------+

================================================================================================
=====================================================================
--- Time Elapsed: 50.608113ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=====================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Black
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
---------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Black
  Configuration State Health              :  Black
    Dev State                             :  provisioned
    App State                             :  cfg refresh error
  Operational State Health                :  Black
    Cluster Health                        :  Black
      Operational State                   :  true
      Peer Operational State              :  false
      Peer Keepalive Operational State    :  true
    Physical Topology Device Health       :  Red
    Device physical topology errors
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+--------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
```

```
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+-----------------+-----------------------
+---------------------+-----------------------+---------------------
+-----------------------------+---------------+
| 10.20.246.6    | Leaf             |                 |                       |
10.20.246.5         | Leaf                  |
|                            | missing-links |
+---------------+-----------------+-----------------+-----------------------
+---------------------+-----------------------+---------------------
+-----------------------------+---------------+
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+-----------------------+--------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+---------------+--------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+-----------------------+--------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+---------------+--------------------------+
| 10.20.246.6      | 10.20.246.5           | 172.31.254.55           | 10.20.20.15 |
65001             | 65001                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
+------------------+-----------------------+--------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+---------------+--------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.5 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Cluster Health                     :  Black
      Operational State                :  true
      Peer Operational State           :  false
      Peer Keepalive Operational State  :  true
    Physical Topology Device Health    :  Red
    Device physical topology errors
+---------------+-----------------+-----------------+-----------------------
+---------------------+-----------------------+---------------------
+-----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+-----------------+-----------------------
+---------------------+-----------------------+---------------------
+-----------------------------+---------------+
| 10.20.246.5    | Leaf             |                 |                       |
10.20.246.6         | Leaf                  |
|                            | missing-links |
+---------------+-----------------+-----------------+-----------------------
+---------------------+-----------------------+---------------------
+-----------------------------+---------------+
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+-----------------------+--------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+---------------+--------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
```

```
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| 10.20.246.5      | 10.20.246.6          | 172.31.254.156          | 10.20.20.14 |
65001             | 65001                | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Cluster Health                     :  Black
      Operational State                :  true
      Peer Operational State           :  false
      Peer Keepalive Operational State  :  true
    Physical Topology Device Health    :  Red
    Device physical topology errors
+---------------+-----------------+----------------+----------------------
+--------------------+---------------------+---------------------
+----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+----------------+----------------------
+--------------------+---------------------+---------------------
+----------------------------+---------------+
| 10.20.246.4    | Leaf             |                 |                      |
10.20.246.3         | Leaf                 |
|                           | missing-links |
+---------------+-----------------+----------------+----------------------
+--------------------+---------------------+---------------------
+----------------------------+---------------+
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
```

```
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.4      | 10.20.246.3           | 172.31.254.2            | 10.20.20.16 |
65002             | 65002                 | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Cluster Health                     :  Black
      Operational State                :  true
      Peer Operational State           :  false
      Peer Keepalive Operational State  :  true
    Physical Topology Device Health    :  Red
    Device physical topology errors
+----------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+----------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+----------------+
| 10.20.246.3    | Leaf             |                 |                      |
10.20.246.4         | Leaf                 |
|                           | missing-links |
+----------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+----------------+
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.3      | 10.20.246.4           | 172.31.254.246          | 10.20.20.17 |
65002             | 65002                 | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
------------------------------------------------------------------------------------------
------------------------------------------------------------------------

==========================================================================================
========================================================================
--- Time Elapsed: 52.950728ms ---
```

Choose Not to Shut Down the MCT Leaf Nodes

The following sample output configures no shut down for the MCT leaf nodes on interface port channel:

```
NHF-Leaf1(config)# interface Port-channel 64
NHF-Leaf1(config-Port-channel-64)# no shutdown
NHF-Leaf1(config-Port-channel-64)#

NH-Leaf1(config)# interface Port-channel 64
NH-Leaf1(config-Port-channel-64)#no  shutdown
NH-Leaf1(config-Port-channel-64)#


Need to wait for 2 min to get all updates and fabric comes back to healthy

(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE
| DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine
| provisioned  | cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine
| provisioned  | cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf
| provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf
| provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf
| provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf
| provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+


(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================
Fabric Name                          :  fab3
Fabric Type                          :  clos
Fabric Health                        :  Green
 Fabric Status                       :  configure-success
 Fabric Level Physical Topology Health  :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================
--- Time Elapsed: 36.802214ms ---
```

```
(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
============================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Green
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ----------------------------------------------------------
 Device IP [Role]                         :  10.20.246.5 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ----------------------------------------------------------
 Device IP [Role]                         :  10.20.246.1 [Spine]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ----------------------------------------------------------
 Device IP [Role]                         :  10.20.246.2 [Spine]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ----------------------------------------------------------
 Device IP [Role]                         :  10.20.246.4 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ----------------------------------------------------------
 Device IP [Role]                         :  10.20.246.3 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
```

```
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
------------------------------------------------------------

============================================================
--- Time Elapsed: 43.465397ms ---
```

Remove Cluster Config from one MCT Leaf

The following sample output removes cluster configuration from one MCT leaf:

```
NHF-Leaf1(config)# do show running-config cluster
cluster fab3-cluster-1
 peer 10.20.20.14
 peer-interface Port-channel 64
 peer-keepalive
  auto
 !
 member vlan all
 member bridge-domain all
!
NHF-Leaf1(config)# no cluster



(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+--------------+---------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |   APP STATE   | CONFIG
GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+---------------
+-------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync   |
NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync   |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg refreshed |
MD                | MCT-C,MCT-PA    | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+---------------
+-------------------+-----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
```

```
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 80.76357ms ---

(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
=============================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+----------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+----------------+
| 10.20.246.6 | Leaf  | Green               | Black             | Black         |
| 10.20.246.5 | Leaf  | Red                 | Black             | Black         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+----------------+

================================================================================================
=============================================
--- Time Elapsed: 39.976662ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=============================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Black
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
---------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Black
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Black
    Cluster Health                        :  Black
      Operational State                   :  true
      Peer Operational State              :  false
      Peer Keepalive Operational State    :  false
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ------------------------------------------------------
 Device IP [Role]                         :  10.20.246.5 [Leaf]
  Device Health                           :  critical
  Configuration State Health              :  Red
    Dev State                             :  provisioned
    App State                             :  cfg refreshed
```

```
  Operational State Health              :  Black
    Cluster Health                      :  Black
      Operational State                 :  false
      Peer Operational State            :  false
      Peer Keepalive Operational State  :  false
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.1 [Spine]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.2 [Spine]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.4 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.3 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
-------------------------------------------------------------------------------------------------
-----------------------------------------------

=================================================================================================
==============================================
--- Time Elapsed: 38.719176ms ---
```

Trigger DRC to Reconcile Cluster Config

The following sample output initiates DRC and reconciles cluster configuration:

```
efa fabric debug device drift --device-ip 10.20.246.5 --name fab3 --reconcile


Fabric Service Response:
Config Drift: Global Config
+--------------------------------+-------------+----------------+
|            CONFIG              |  APP STATE  | EXPECTED VALUE |
+--------------------------------+-------------+----------------+
```

```
| Mtu                           | cfg-in-sync | 9216           |
| IPMtu                         | cfg-in-sync | 9100           |
| AnycastMac                    | cfg-in-sync | 0201.0101.0101 |
| IPV6AnycastMac                | cfg-in-sync | 0201.0101.0102 |
| MacAgingConversationalTimeout | cfg-in-sync | 300            |
| MacAgingTimeout               | cfg-in-sync | 1800           |
| MacMoveLimit                  | cfg-in-sync | 20             |
| MacMoveDetect                 | cfg-in-sync | true           |
+-------------------------------+-------------+----------------+
Config Drift: EVPN
+------+-------------+---------------------------+
| NAME |  APP STATE  |       CHILD CONFIG        |
+------+-------------+---------------------------+
| fab3 | cfg-in-sync | SwEvpnName                |
| fab3 | cfg-in-sync | DuplicateMacTimerMaxCount |
| fab3 | cfg-in-sync | DuplicateMacTimer         |
| fab3 | cfg-in-sync | RouteTarget               |
| fab3 | cfg-in-sync | Rd                        |
+------+-------------+---------------------------+
Config Drift: Overlay Gateway
+------+-------------+------------------------+
| NAME |  APP STATE  |      CHILD CONFIG      |
+------+-------------+------------------------+
| fab3 | cfg-in-sync | SwOverlayGwName        |
| fab3 | cfg-in-sync | VtepLoopbackPortNumber |
| fab3 | cfg-in-sync | MapVniAuto             |
| fab3 | cfg-in-sync | Activate               |
+------+-------------+------------------------+
Config Drift: Cluster
+---------+---------------+--------------------------------+
|  NAME   |   APP STATE   |          CHILD CONFIG          |
+---------+---------------+--------------------------------+
| Cluster | cfg-refreshed | ClusterName                    |
| Cluster | cfg-refreshed | MCTPeerName::0:Port-channel:64 |
| Cluster | cfg-refreshed | ClusterKeepaliveAuto::0        |
+---------+---------------+--------------------------------+
Config Drift: Interface
+------+-------------+--------------+-----------------------------------+
| NAME |  APP STATE  |   INT TYPE   |           CHILD CONFIG            |
+------+-------------+--------------+-----------------------------------+
| 0/54 | cfg-in-sync | ethernet     | IP:0/54:ethernet:10.10.10.33/31   |
| 0/54 | cfg-in-sync | ethernet     | IPPimSparse:0/54:ethernet:false   |
| 0/54 | cfg-in-sync | ethernet     | BFD:0/54:ethernet:3:300:300       |
| 0/52 | cfg-in-sync | ethernet     | IP:0/52:ethernet:10.10.10.35/31   |
| 0/52 | cfg-in-sync | ethernet     | IPPimSparse:0/52:ethernet:false   |
| 0/52 | cfg-in-sync | ethernet     | BFD:0/52:ethernet:3:300:300       |
| 1    | cfg-in-sync | loopback     | IP:1:loopback:172.31.254.156/32   |
| 2    | cfg-in-sync | loopback     | IP:2:loopback:172.31.254.210/32   |
| 64   | cfg-in-sync | port-channel | IP:64:port-channel:10.20.20.15/31 |
+------+-------------+--------------+-----------------------------------+
Config Drift: Router BGP
+--------+-------------+-----------------------------------+
|  TYPE  |  APP STATE  |           CHILD CONFIG            |
+--------+-------------+-----------------------------------+
| Global | cfg-in-sync | BgpDynamicPeerListenLimit         |
| Global | cfg-in-sync | PeerGroupInfo                     |
| Global | cfg-in-sync | BgpNeighbor                       |
| Global | cfg-in-sync | BgpMCTBFDNeighbor                 |
| Global | cfg-in-sync | BgpMCTNeighbor                    |
| Global | cfg-in-sync | RouterID                          |
| Global | cfg-in-sync | LocalAsn                          |
| Global | cfg-in-sync | FastExternalFallOver              |
| Global | cfg-in-sync | CapabilityAs4Enable               |
| Global | cfg-in-sync | BfdMultiplier                     |
```

```
| Global | cfg-in-sync | BfdTx                                |
| Global | cfg-in-sync | BfdRx                                |
| Global | cfg-in-sync | BgpIPV4Network                       |
| Global | cfg-in-sync | BgpIPV4NetworkGracefulRestart        |
| Global | cfg-in-sync | BgpL2EVPNNetworkGracefulRestart      |
| Global | cfg-in-sync | BgpL2EVPNNetworkEnablePeerAsCheck    |
| Global | cfg-in-sync | BgpL2EVPNNetworkEncapsulation        |
| Global | cfg-in-sync | BgpL2EVPNNetworkNextHopUnchanged     |
| Global | cfg-in-sync | BgpL2EVPNNetworkActivate             |
| Global | cfg-in-sync | BgpIPV4NetworkMaxPath                |
+--------+-------------+--------------------------------------+
+-------------+---------+-------+
| CONFIG TYPE | STATUS  | ERROR |
+-------------+---------+-------+
| MCT         | Success |       |
+-------------+---------+-------+



Wait for 1 minute to get updated status

(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details
```

```
--- Time Elapsed: 65.918511ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================================
================================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+----------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+----------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green          |
| 10.20.246.5 | Leaf  | Green               | Green             | Green          |
| 10.20.246.1 | Spine | Green               | Green             | Green          |
| 10.20.246.2 | Spine | Green               | Green             | Green          |
| 10.20.246.4 | Leaf  | Green               | Green             | Green          |
| 10.20.246.3 | Leaf  | Green               | Green             | Green          |
+-------------+-------+---------------------+-------------------+----------------+

================================================================================================================
================================================================================
--- Time Elapsed: 37.887051ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================================
================================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Green
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
----------------------------------------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ------------------------------------------------------
 Device IP [Role]                         :  10.20.246.5 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ------------------------------------------------------
 Device IP [Role]                         :  10.20.246.1 [Spine]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
```

```
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.2 [Spine]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.4 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.3 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------

===========================================================================================
=================================================================
--- Time Elapsed: 37.954713ms ---
```

Remove Leaf to Spine Links between Devices

The following sample output removes links between leaf and spine devices:

```
10.20.246.5,10.20.246.6  to 10.20.246.2 spine  links removed
10.20.246.3,10.20.246.4 to 10.20.246.1  spine links removed


(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+-----------------------------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |     APP STATE     |
CONFIG GEN REASON |         PENDING CONFIGS          | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+-----------------------------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg refresh error |
LD,IU             | BGP-U,BGP-D,INTIP-D               | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg refresh error |
LD,IU             | BGP-U,BGP-D,INTIP-U,INTIP-D       | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg refresh error |
LD,IU             | BGP-U,BGP-D,INTIP-U,INTIP-D       | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg refresh error |
```

```
LD,IU             | BGP-U,BGP-D,INTIP-D              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg refresh error |
LD,IU             | BGP-U,BGP-D,INTIP-D              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg refresh error |
LD,IU             | MCT-C,MCT-PA,BGP-U,BGP-D,INTIP-D | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------------
+------------------+-----------------------------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 166.73422ms ---




(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+----------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+----------------+
| 10.20.246.6 | Leaf  | Black               | Black             | Black         |
| 10.20.246.5 | Leaf  | Black               | Black             | Black         |
| 10.20.246.1 | Spine | Black               | Black             | Black         |
| 10.20.246.2 | Spine | Black               | Black             | Black         |
| 10.20.246.4 | Leaf  | Black               | Black             | Black         |
| 10.20.246.3 | Leaf  | Black               | Black             | Black         |
+-------------+-------+---------------------+-------------------+----------------+

================================================================================================
=====================================================================
--- Time Elapsed: 54.258428ms ---


(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
```

```
 Fabric Status                          :  configure-success
 Fabric Level Physical Topology Health  :  Green
-------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                       :  10.20.246.6 [Leaf]
  Device Health                         :  Black
  Configuration State Health            :  Black
    Dev State                           :  provisioned
    App State                           :  cfg refresh error
  Operational State Health              :  Black
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Red
    Device physical topology errors
+----------------+------------------+-----------------+----------------------
+---------------------+----------------------+----------------------
+-----------------------------+----------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+------------------+-----------------+----------------------
+---------------------+----------------------+----------------------
+-----------------------------+----------------+
| 10.20.246.6    | Leaf             |                 |                      |
10.20.246.2         | Spine                 |
|                           | missing-links |
+----------------+------------------+-----------------+----------------------
+---------------------+----------------------+----------------------
+-----------------------------+----------------+
    Underlay Topology Device Health     :  Red
    Device underlay topology errors
+------------------+----------------------+--------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+--------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| 10.20.246.6      | 10.20.246.2           | 172.31.254.55           | 10.10.10.37 |
65001             | 64512                 | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
| 10.20.246.6      | 10.20.246.2           | 172.31.254.55           | 10.10.10.37 |
65001             | 64512                 | default-vrf | l2vpn              |
evpn          | CONN           | session_not_established |
+------------------+----------------------+--------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
 -------------------------------------------------------
 Device IP [Role]                       :  10.20.246.5 [Leaf]
  Device Health                         :  Black
  Configuration State Health            :  Black
    Dev State                           :  provisioned
    App State                           :  cfg refresh error
  Operational State Health              :  Black
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Red
    Device physical topology errors
+----------------+------------------+-----------------+----------------------
+---------------------+----------------------+----------------------
+-----------------------------+----------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
```

```
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
| 10.20.246.5    | Leaf            |                |                       |
10.20.246.2         | Spine                |
|                            | missing-links |
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.5      | 10.20.246.2           | 172.31.254.156          | 10.10.10.34 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.5      | 10.20.246.2           | 172.31.254.156          | 10.10.10.34 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Physical Topology Device Health    :  Red
    Device physical topology errors
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
| 10.20.246.1    | Spine           |                |                       |
10.20.246.4         | Leaf                 |
|                            | missing-links |
| 10.20.246.1    | Spine           |                |                       |
10.20.246.3         | Leaf                 |
|                            | missing-links |
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
```

```
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+-----------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.1      | 10.20.246.4           | 172.31.254.144          | 10.10.10.40 |
64512             | 65002                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.1      | 10.20.246.4           | 172.31.254.144          | 10.10.10.40 |
64512             | 65002                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
| 10.20.246.1      | 10.20.246.3           | 172.31.254.144          | 10.10.10.45 |
64512             | 65002                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.1      | 10.20.246.3           | 172.31.254.144          | 10.10.10.45 |
64512             | 65002                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+-----------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Physical Topology Device Health    :  Red
    Device physical topology errors
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
| 10.20.246.2    | Spine            |                 |                       |
10.20.246.6         | Leaf                 |
|                            | missing-links |
| 10.20.246.2    | Spine            |                 |                       |
10.20.246.5         | Leaf                 |
|                            | missing-links |
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+-----------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+-----------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.2      | 10.20.246.5           | 172.31.254.205          | 10.10.10.35 |
64512             | 65001                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.2      | 10.20.246.5           | 172.31.254.205          | 10.10.10.35 |
64512             | 65001                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
```

```
| 10.20.246.2      | 10.20.246.6           | 172.31.254.205          | 10.10.10.36 |
64512             | 65001                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.2      | 10.20.246.6           | 172.31.254.205          | 10.10.10.36 |
64512             | 65001                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Red
    Device physical topology errors
+---------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+-----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+-----------------------------+---------------+
| 10.20.246.4    | Leaf             |                 |                      |
10.20.246.1         | Spine                 |
|                           | missing-links |
+---------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+-----------------------------+---------------+
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.4      | 10.20.246.1           | 172.31.254.2            | 10.10.10.41 |
65002             | 64512                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.4      | 10.20.246.1           | 172.31.254.2            | 10.10.10.41 |
65002             | 64512                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Red
    Device physical topology errors
```

```
+----------------+-----------------+----------------+----------------------
+--------------------+--------------------+---------------------
+---------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+-----------------+----------------+----------------------
+--------------------+--------------------+---------------------
+---------------------------+---------------+
| 10.20.246.3    | Leaf             |                 |                       |
10.20.246.1         | Spine                |
|                             | missing-links |
+----------------+-----------------+----------------+----------------------
+--------------------+--------------------+---------------------
+---------------------------+---------------+
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+-------------------+----------------------+-------------+-------------------
+---------------+----------------+------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+-------------------+----------------------+-------------+-------------------
+---------------+----------------+------------------------+
| 10.20.246.3      | 10.20.246.1           | 172.31.254.246          | 10.10.10.44 |
65002             | 64512                  | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
| 10.20.246.3      | 10.20.246.1           | 172.31.254.246          | 10.10.10.44 |
65002             | 64512                  | default-vrf | l2vpn              |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+-------------------+----------------------+-------------+-------------------
+---------------+----------------+------------------------+
-------------------------------------------------------------------------------------------
---------------------------------------------------------------------

===========================================================================================
=====================================================================
--- Time Elapsed: 71.492989ms ---
```

10.20.246.5,10.20.246.6 leaf to 10.20.246.2 spine links added. 10.20.246.3,10.20.246.4 leaf to 10.20.246.1 spine links added.

```
(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+----------------+--------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+----------------+--------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA               | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA               | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA               | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA               | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
```

```
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 63.575222ms ---

(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
================================================================
--- Time Elapsed: 35.046579ms ---


(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
------------------------------------------------------------------------------------------------
----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.6 [Leaf]
```

```
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.5 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.1 [Spine]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.2 [Spine]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.4 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.3 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
-------------------------------------------------------------------------------------------------
-------------------------------------------------------------------

=================================================================================================
=================================================================
--- Time Elapsed: 39.076563ms ---
```

Add Incorrect Links between Two Spines

The following sample output configures invalid links between spine devices:

```
NH-1(config)# interface Ethernet 0/13-15
NH-1(conf-if-eth-0/13-15)# no shutdown


(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+----------------+--------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |     APP STATE     |
CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+----------------+--------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg refresh error |
LA                | INTIP-C         | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg refresh error |
LA                | INTIP-C,INTIP-U | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync       |
NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync       |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync       |
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync       |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------------
+-------------------+----------------+--------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 83.368638ms ---


(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
======================================================================
Fabric Name                           :  fab3
Fabric Type                           :  clos
Fabric Health                         :  Black
 Fabric Status                        :  configure-success
 Fabric Level Physical Topology Health  :  Green
Fabric Device Health
+-------------+-------+--------------------+------------------+---------------+
```

```
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Black               | Red               | Black         |
| 10.20.246.2 | Spine | Black               | Red               | Black         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
=====================================================================
--- Time Elapsed: 37.140239ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
------------------------------------------------------------------------------------------------
------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.6 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.5 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.1 [Spine]
  Device Health                          :  Black
  Configuration State Health             :  Black
    Dev State                            :  provisioned
    App State                            :  cfg refresh error
  Operational State Health               :  Red
    Physical Topology Device Health      :  Red
    Device physical topology errors
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+----------------------------+-----------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |      ERROR      |
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
```

```
+---------------------------+-----------------+
| 10.20.246.1    | Spine            |                 | 0/14                 |
10.20.246.2         | Spine                 |                      |
0/14                       | incorrect-links |
| 10.20.246.1    | Spine            |                 | 0/15                 |
10.20.246.2         | Spine                 |                      |
0/15                       | incorrect-links |
| 10.20.246.1    | Spine            |                 | 0/13                 |
10.20.246.2         | Spine                 |                      |
0/13                       | incorrect-links |
+----------------+-----------------+-----------------+-----------------------
+--------------------+----------------------+---------------------
+---------------------------+-----------------+
    Underlay Topology Device Health     :  Green
 ---------------------------------------------------------
 Device IP [Role]                       :  10.20.246.2 [Spine]
  Device Health                         :  Black
  Configuration State Health            :  Black
    Dev State                           :  provisioned
    App State                           :  cfg refresh error
  Operational State Health              :  Red
    Physical Topology Device Health     :  Red
    Device physical topology errors
+----------------+-----------------+-----------------+-----------------------
+--------------------+----------------------+---------------------
+---------------------------+-----------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |      ERROR      |
+----------------+-----------------+-----------------+-----------------------
+--------------------+----------------------+---------------------
+---------------------------+-----------------+
| 10.20.246.2    | Spine            |                 | 0/14                 |
10.20.246.1         | Spine                 |                      |
0/14                       | incorrect-links |
| 10.20.246.2    | Spine            |                 | 0/15                 |
10.20.246.1         | Spine                 |                      |
0/15                       | incorrect-links |
| 10.20.246.2    | Spine            |                 | 0/13                 |
10.20.246.1         | Spine                 |                      |
0/13                       | incorrect-links |
+----------------+-----------------+-----------------+-----------------------
+--------------------+----------------------+---------------------
+---------------------------+-----------------+
    Underlay Topology Device Health     :  Green
 ---------------------------------------------------------
 Device IP [Role]                       :  10.20.246.4 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ---------------------------------------------------------
 Device IP [Role]                       :  10.20.246.3 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
```

```
    Underlay Topology Device Health     :  Green
------------------------------------------------------------------------------------------
---------------------------------------------------------------------

==========================================================================================
=====================================================================
--- Time Elapsed: 51.504456ms ---
```

Remove Incorrect Links between Two Spines

The following sample output removes invalid links between spine devices:

```
NH-1(conf-if-eth-0/13-15)# shutdown
As there is spine to spine deleted config the state is in cfg refreshed, configure fabric
Or DRC trigger will move it to healthy

(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Red
+-------------+-----+-----------+-------+-------+-------------+---------------
+------------------+----------------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |   APP STATE   | CONFIG
GEN REASON |    PENDING CONFIGS    | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+-------------+---------------
+------------------+----------------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg refreshed |
LA,LD,IU          | BGP-U,INTIP-C         | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg refreshed |
LA,LD,IU          | BGP-U,INTIP-C,INTIP-U | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                | NA                    | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                | NA                    | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync   |
NA                | NA                    | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync   |
NA                | NA                    | 2       | 1     |
+-------------+-----+-----------+-------+-------+-------------+---------------
+------------------+----------------------+---------+-------+
CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 82.219266ms ---

(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
=====================================================================
Fabric Name                          :  fab3
```

```
Fabric Type                              :  clos
Fabric Health                            :  Red
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Red                 | healthy           | Red           |
| 10.20.246.2 | Spine | Red                 | healthy           | Red           |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
=================================================================
(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Red
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ---------------------------------------------------------
 Device IP [Role]                         :  10.20.246.5 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ---------------------------------------------------------
 Device IP [Role]                         :  10.20.246.1 [Spine]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  Green
    App State                             :  cfg refreshed
  Operational State Health                :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ---------------------------------------------------------
 Device IP [Role]                         :  10.20.246.2 [Spine]
  Device Health                           :  Red
  Configuration State Health              :  Red
    Dev State                             :  provisioned
    App State                             :  cfg refreshed
```

```
  Operational State Health             :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
------------------------------------------------------------------------------------------
---------------------------------------------------------------------

==========================================================================================
=====================================================================
--- Time Elapsed: 35.201349ms ---
```

Post Configure Fabric or DRC

The following is an example of sample output for post fabric or DRC configuration:

```
(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+--------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+--------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA               | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA               | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+--------+-------+

(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
=====================================================================
Fabric Name                            :  fab3
Fabric Type                            :  clos
```

```
Fabric Health                           :  Green
 Fabric Status                          :  configure-success
 Fabric Level Physical Topology Health  :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
=================================================================
--- Time Elapsed: 37.545522ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=================================================================
Fabric Name                            :  fab3
Fabric Type                            :  clos
Fabric Health                          :  Green
 Fabric Status                         :  configure-success
 Fabric Level Physical Topology Health  :  Green
------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                      :  10.20.246.6 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.5 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
```

```
  Operational State Health              :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.4 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.3 [Leaf]
  Device Health                         :  Green
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Green
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Green
-------------------------------------------------------------------------------------------------
-----------------------------------------------------------------------

=================================================================================================
=====================================================================
--- Time Elapsed: 46.001168ms ---
```

Delete Router BGP Configuration in a Spine Device

The following sample output deletes router BGP configuration from a spine device:

```
Welcome to the Extreme SLX-OS Software
admin connected from 134.141.25.99 using ssh on NH-1
NH-1# conf t
Entering configuration mode terminal
NH-1(config)# no router bgp
NH-1(config)#

(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+--------------+---------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |   APP STATE   | CONFIG
GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+---------------
+-------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg refreshed |
ASN               | BGP-U           | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync   |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync   |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+---------------
```

```
+------------------+----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 82.973688ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
=================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Black             | Black         |
| 10.20.246.5 | Leaf  | Green               | Black             | Black         |
| 10.20.246.1 | Spine | Red                 | Black             | Black         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Black             | Black         |
| 10.20.246.3 | Leaf  | Green               | Black             | Black         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
=================================================================
--- Time Elapsed: 46.626729ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Black
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Black
```

```
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Black
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.6      | 10.20.246.1           | 172.31.254.55           | 10.10.10.39 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.6      | 10.20.246.1           | 172.31.254.55           | 10.10.10.39 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.5 [Leaf]
  Device Health                         :  Black
  Configuration State Health            :  Green
    Dev State                           :  provisioned
    App State                           :  cfg in-sync
  Operational State Health              :  Black
    Cluster Health                      :  Green
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.5      | 10.20.246.1           | 172.31.254.156          | 10.10.10.32 |
65001             | 64512                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.5      | 10.20.246.1           | 172.31.254.156          | 10.10.10.32 |
65001             | 64512                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                       :  10.20.246.1 [Spine]
  Device Health                         :  Black
  Configuration State Health            :  Red
    Dev State                           :  provisioned
    App State                           :  cfg refreshed
  Operational State Health              :  Black
    Physical Topology Device Health     :  Green
    Underlay Topology Device Health     :  Black
```

```
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.1      | 10.20.246.5           | 172.31.254.144          | 10.10.10.33 |
64512             | 65001                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.5           | 172.31.254.144          | 10.10.10.33 |
64512             | 65001                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.6           | 172.31.254.144          | 10.10.10.38 |
64512             | 65001                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.6           | 172.31.254.144          | 10.10.10.38 |
64512             | 65001                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.4           | 172.31.254.144          | 10.10.10.40 |
64512             | 65002                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.4           | 172.31.254.144          | 10.10.10.40 |
64512             | 65002                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.3           | 172.31.254.144          | 10.10.10.45 |
64512             | 65002                 | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.1      | 10.20.246.3           | 172.31.254.144          | 10.10.10.45 |
64512             | 65002                 | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Black
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
```

```
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.4      | 10.20.246.1          | 172.31.254.2            | 10.10.10.41 |
65002             | 64512                | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.4      | 10.20.246.1          | 172.31.254.2            | 10.10.10.41 |
65002             | 64512                | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Black
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.3      | 10.20.246.1          | 172.31.254.246          | 10.10.10.44 |
65002             | 64512                | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.3      | 10.20.246.1          | 172.31.254.246          | 10.10.10.44 |
65002             | 64512                | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
-------------------------------------------------------------------------------------------------
-------------------------------------------------------------------------

=================================================================================================
=========================================================================
--- Time Elapsed: 59.059441ms ---
```

Trigger DRC to Reconcile Router BGP Config on Switch

The following sample output initiates DRC to reconcile router BGP configuration on a switch:

```
(efa:user)user@dev-server:~$ efa fabric debug device drift --device-ip 10.20.246.1 --name
fab3 --reconcile


Fabric Service Response:
Config Drift: Global Config
+--------+-------------+----------------+
| CONFIG |  APP STATE  | EXPECTED VALUE |
+--------+-------------+----------------+
| Mtu    | cfg-in-sync | 9216           |
| IPMtu  | cfg-in-sync | 9100           |
+--------+-------------+----------------+
```

```
Config Drift: EVPN
+------+-----------+--------------+
| NAME | APP STATE | CHILD CONFIG |
+------+-----------+--------------+
+------+-----------+--------------+
Config Drift: Overlay Gateway
+------+-----------+--------------+
| NAME | APP STATE | CHILD CONFIG |
+------+-----------+--------------+
+------+-----------+--------------+
Config Drift: Cluster
+------+-----------+--------------+
| NAME | APP STATE | CHILD CONFIG |
+------+-----------+--------------+
+------+-----------+--------------+
Config Drift: Interface
+------+-------------+----------+---------------------------------+
| NAME |  APP STATE  | INT TYPE |          CHILD CONFIG           |
+------+-------------+----------+---------------------------------+
| 0/31 | cfg-in-sync | ethernet | IP:0/31:ethernet:10.10.10.41/31 |
| 0/31 | cfg-in-sync | ethernet | IPPimSparse:0/31:ethernet:false |
| 0/31 | cfg-in-sync | ethernet | BFD:0/31:ethernet:3:300:300     |
| 0/21 | cfg-in-sync | ethernet | IP:0/21:ethernet:10.10.10.32/31 |
| 0/21 | cfg-in-sync | ethernet | IPPimSparse:0/21:ethernet:false |
| 0/21 | cfg-in-sync | ethernet | BFD:0/21:ethernet:3:300:300     |
| 0/24 | cfg-in-sync | ethernet | IP:0/24:ethernet:10.10.10.39/31 |
| 0/24 | cfg-in-sync | ethernet | IPPimSparse:0/24:ethernet:false |
| 0/24 | cfg-in-sync | ethernet | BFD:0/24:ethernet:3:300:300     |
| 0/32 | cfg-in-sync | ethernet | IP:0/32:ethernet:10.10.10.44/31 |
| 0/32 | cfg-in-sync | ethernet | IPPimSparse:0/32:ethernet:false |
| 0/32 | cfg-in-sync | ethernet | BFD:0/32:ethernet:3:300:300     |
| 1    | cfg-in-sync | loopback | IP:1:loopback:172.31.254.144/32 |
+------+-------------+----------+---------------------------------+
Config Drift: Router BGP
+--------+---------------+-----------------------------------+
|  TYPE  |   APP STATE   |           CHILD CONFIG            |
+--------+---------------+-----------------------------------+
| Global | cfg-refreshed | BgpDynamicPeerListenLimit         |
| Global | cfg-refreshed | PeerGroupInfo                     |
| Global | cfg-refreshed | BgpNeighbor                       |
| Global | cfg-in-sync   | BgpMCTBFDNeighbor                 |
| Global | cfg-in-sync   | BgpMCTNeighbor                    |
| Global | cfg-refreshed | RouterID                          |
| Global | cfg-refreshed | LocalAsn                          |
| Global | cfg-refreshed | FastExternalFallOver              |
| Global | cfg-refreshed | CapabilityAs4Enable               |
| Global | cfg-in-sync   | BfdMultiplier                     |
| Global | cfg-in-sync   | BfdTx                             |
| Global | cfg-in-sync   | BfdRx                             |
| Global | cfg-in-sync   | BgpIPV4Network                    |
| Global | cfg-refreshed | BgpIPV4NetworkGracefulRestart     |
| Global | cfg-refreshed | BgpL2EVPNNetworkGracefulRestart   |
| Global | cfg-refreshed | BgpL2EVPNRetainRtAll              |
| Global | cfg-refreshed | BgpL2EVPNNetworkEnablePeerAsCheck |
| Global | cfg-refreshed | BgpL2EVPNNetworkEncapsulation     |
| Global | cfg-refreshed | BgpL2EVPNNetworkNextHopUnchanged  |
| Global | cfg-refreshed | BgpL2EVPNNetworkActivate          |
| Global | cfg-refreshed | BgpIPV4NetworkMaxPath             |
+--------+---------------+-----------------------------------+
+-------------+---------+-------+
| CONFIG TYPE | STATUS  | ERROR |
+-------------+---------+-------+
| routerbgp   | Success |       |
+-------------+---------+-------+
```

```
--- Time Elapsed: 37.210932444s ---




(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+--------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+--------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA     | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA     | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2      | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2      | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2      | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2      | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+--------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 65.881523ms ---


(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
=====================================================================
Fabric Name                            :  fab3
Fabric Type                            :  clos
Fabric Health                          :  Green
 Fabric Status                         :  configure-success
 Fabric Level Physical Topology Health  :  Green
Fabric Device Health
+-------------+-------+--------------------+------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+--------------------+------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
```

```
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
================================================================
--- Time Elapsed: 41.067435ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
------------------------------------------------------------------------------------------------
------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.6 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ---------------------------------------------------------
 Device IP [Role]                        :  10.20.246.5 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ---------------------------------------------------------
 Device IP [Role]                        :  10.20.246.1 [Spine]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ---------------------------------------------------------
 Device IP [Role]                        :  10.20.246.2 [Spine]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ---------------------------------------------------------
 Device IP [Role]                        :  10.20.246.4 [Leaf]
```

```
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.3 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
-------------------------------------------------------------------------------------------------
-------------------------------------------------------------------

=================================================================================================
=====================================================================
--- Time Elapsed: 32.967631ms ---
```

Delete BGP Neighbors from Leaf Device

The following sample output deletes BGP neighbors from a leaf device:

```
10.20.246.4
NH-Leaf2# show bgp evpn summary
  BGP4 Summary
  Router ID: 172.31.254.2   Local AS Number: 65002
  Confederation Identifier: not configured
  Confederation Peers:
  Maximum Number of IP ECMP Paths Supported for Load Sharing: 1
  Number of Neighbors Configured: 2, UP: 2
  Number of Routes Installed: 5, Uses 830 bytes
  Number of Routes Advertising to All Neighbors: 2 (1 entries), Uses 76 bytes
  Number of Attribute Entries Installed: 4, Uses 764 bytes
  d: Dynamically created based on a listen range command
  Dynamically created neighbors: 0/100(max)
  A: Auto Discovered Neighbors using LLDP
  Auto Neighbors Count: 0
  '+': Data in InQueue '>': Data in OutQueue '-': Clearing
  '*': Update Policy 'c': Group change 'p': Group change Pending
  'r': Restarting 's': Stale '^': Up before Restart '<': EOR waiting
  '$': Learning-Phase (for Delayed Route Calculation)
  '#': RIB-in Phase
  'D': Dampening enabled
  Neighbor Address  AS#         State     Time     Rt:Accepted Filtered Sent     ToSend
  10.10.10.41       64512       ESTAB   0h4m40s       2        0        1        0
  10.10.10.43       64512       ESTAB   2h56m9s       2        0        1        0
NH-Leaf2# conf t
Entering configuration mode terminal
NH-Leaf2(config)# router bgp
NH-Leaf2(config-bgp-router)# do show running-config router bgp
router bgp
 local-as 65002
 capability as4-enable
 fast-external-fallover
 neighbor spine-group peer-group
 neighbor spine-group remote-as 64512
 neighbor spine-group description To Spine
```

```
 neighbor spine-group bfd
 neighbor 10.10.10.41 peer-group spine-group
 neighbor 10.10.10.43 peer-group spine-group
 neighbor 10.20.20.16 remote-as 65002
 neighbor 10.20.20.16 next-hop-self
 neighbor 10.20.20.16 bfd
 address-family ipv4 unicast
  network 172.31.254.110/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
  neighbor spine-group encapsulation vxlan
  neighbor spine-group next-hop-unchanged
  neighbor spine-group enable-peer-as-check
  neighbor spine-group activate
 !
!
NH-Leaf2(config-bgp-router)# no  neighbor 10.10.10.41 peer-group spine-group
%Warning: Clean up BGP routes for peer

10.20.246.5
NHF-Leaf1# show bgp evpn summary
  BGP4 Summary
  Router ID: 172.31.254.156   Local AS Number: 65001
  Confederation Identifier: not configured
  Confederation Peers:
  Maximum Number of IP ECMP Paths Supported for Load Sharing: 1
  Number of Neighbors Configured: 2, UP: 2
  Number of Routes Installed: 5, Uses 830 bytes
  Number of Routes Advertising to All Neighbors: 2 (1 entries), Uses 76 bytes
  Number of Attribute Entries Installed: 4, Uses 764 bytes
  d: Dynamically created based on a listen range command
  Dynamically created neighbors: 0/100(max)
  A: Auto Discovered Neighbors using LLDP
  Auto Neighbors Count: 0
  '+': Data in InQueue '>': Data in OutQueue '-': Clearing
  '*': Update Policy 'c': Group change 'p': Group change Pending
  'r': Restarting 's': Stale '^': Up before Restart '<': EOR waiting
  '$': Learning-Phase (for Delayed Route Calculation)
  '#': RIB-in Phase
  'D': Dampening enabled
  Neighbor Address  AS#         State     Time     Rt:Accepted Filtered Sent     ToSend
  10.10.10.32       64512       ESTAB   0h4m31s       2        0        1        0
  10.10.10.34       64512       ESTAB   0h39m28s      2        0        1        0
NHF-Leaf1# conf t
Entering configuration mode terminal
NHF-Leaf1(config)# do show running-config router bgp
router bgp
 local-as 65001
 capability as4-enable
 fast-external-fallover
 neighbor spine-group peer-group
 neighbor spine-group remote-as 64512
 neighbor spine-group description To Spine
 neighbor spine-group bfd
 neighbor 10.10.10.32 peer-group spine-group
 neighbor 10.10.10.34 peer-group spine-group
 neighbor 10.20.20.14 remote-as 65001
 neighbor 10.20.20.14 next-hop-self
 neighbor 10.20.20.14 bfd
```

```
 address-family ipv4 unicast
  network 172.31.254.210/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
  neighbor spine-group encapsulation vxlan
  neighbor spine-group next-hop-unchanged
  neighbor spine-group enable-peer-as-check
  neighbor spine-group activate
 !
!
NHF-Leaf1(config)# router bgp
NHF-Leaf1(config-bgp-router)# no neighbor 10.10.10.32 peer-group spine-group
%Warning: Clean up BGP routes for peer



(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+--------------+---------------
+--------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |   APP STATE   | CONFIG
GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+---------------
+--------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync   |
NA                 | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync   |
NA                 | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync   |
NA                 | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg refreshed |
BGPU               | BGP-C,INTIP-C   | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg refreshed |
BGPU               | BGP-C           | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync   |
NA                 | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+---------------
+--------------------+-----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details
```

```
--- Time Elapsed: 148.017748ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================================
================================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Red                 | Black             | Black         |
| 10.20.246.1 | Spine | Green               | Black             | Black         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Red                 | Black             | Black         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================================
================================================================================
--- Time Elapsed: 38.149955ms ---


(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================================
================================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Black
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
----------------------------------------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
  Device Health                           :  Green
  Configuration State Health              :  Green
    Dev State                             :  provisioned
    App State                             :  cfg in-sync
  Operational State Health                :  Green
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Green
 ------------------------------------------------------
 Device IP [Role]                         :  10.20.246.5 [Leaf]
  Device Health                           :  Black
  Configuration State Health              :  Red
    Dev State                             :  provisioned
    App State                             :  cfg refreshed
  Operational State Health                :  Black
    Cluster Health                        :  Green
    Physical Topology Device Health       :  Green
    Underlay Topology Device Health       :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+-----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
```

```
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| 10.20.246.5      | 10.20.246.1          | 172.31.254.156          | 10.10.10.32 |
65001            | 64512                | default-vrf | l2vpn              |
evpn          |                | neighbor_not_configured |
| 10.20.246.5      | 10.20.246.1          | 172.31.254.156          | 10.10.10.32 |
65001            | 64512                | default-vrf | ipv4               |
unicast       |                | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Black
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Black
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| 10.20.246.1      | 10.20.246.5          | 172.31.254.144          | 10.10.10.33 |
64512            | 65001                | default-vrf | ipv4               |
unicast       | ACTIV          | session_not_established |
| 10.20.246.1      | 10.20.246.5          | 172.31.254.144          | 10.10.10.33 |
64512            | 65001                | default-vrf | l2vpn              |
evpn          | ACTIV          | session_not_established |
| 10.20.246.1      | 10.20.246.4          | 172.31.254.144          | 10.10.10.40 |
64512            | 65002                | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
| 10.20.246.1      | 10.20.246.4          | 172.31.254.144          | 10.10.10.40 |
64512            | 65002                | default-vrf | l2vpn              |
evpn          | CONN           | session_not_established |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Red
    Dev State                          :  provisioned
    App State                          :  cfg refreshed
  Operational State Health             :  Black
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
```

```
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.4      | 10.20.246.1           | 172.31.254.2            | 10.10.10.41 |
65002             | 64512                 | default-vrf | l2vpn              |
evpn          |               | neighbor_not_configured |
| 10.20.246.4      | 10.20.246.1           | 172.31.254.2            | 10.10.10.41 |
65002             | 64512                 | default-vrf | ipv4               |
unicast       |               | neighbor_not_configured |
+------------------+----------------------+-------------------------+-------------
+------------------+----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
-------------------------------------------------------------------------------------------
-------------------------------------------------------------------

===========================================================================================
=========================================================================
--- Time Elapsed: 44.977239ms ---
```

Configure Neighbor again in Switch

The following sample output configures neighbors in a switch:

```
10.20.246.5 to spine 10.20.246.1
NHF-Leaf1(config-bgp-router)# neighbor 10.10.10.32 peer-group spine-group

10.20.246.4
NH-Leaf2(config-bgp-router)# neighbor 10.10.10.41 peer-group spine-group


(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
```

```
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 61.727945ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3
================================================================================================
=================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
Fabric Device Health
+-------------+-------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Green               | Green             | Green         |
| 10.20.246.5 | Leaf  | Green               | Green             | Green         |
| 10.20.246.1 | Spine | Green               | Green             | Green         |
| 10.20.246.2 | Spine | Green               | Green             | Green         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
=================================================================
--- Time Elapsed: 54.912866ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
=================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Green
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
-----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                         :  10.20.246.6 [Leaf]
```

```
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.5 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.1 [Spine]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.2 [Spine]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.4 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.3 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
-------------------------------------------------------------------------------------------------
-------------------------------------------------------------------

=================================================================================================
=====================================================================
--- Time Elapsed: 34.75031ms ---
```

Reload a Leaf Device

The following sample output configures to reload a leaf device:

```
(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Black
+-------------+-----+-----------+-------+-------+-------------+-------------------
+------------------+----------------+--------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |     APP STATE     |
CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+-------------+-------------------
+------------------+----------------+--------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg refresh error |
LD,IU            | BGP-U           | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg refresh error |
LD,IU            | BGP-U           | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync       |
NA               | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync       |
NA               | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg refresh error |
LD               | BGP-U           | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg refresh error |
LD,IU            | BGP-U           | 2       | 1     |
+-------------+-----+-----------+-------+-------+-------------+-------------------
+------------------+----------------+--------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 128.855713ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
=====================================================================
Fabric Name                             :  fab3
Fabric Type                             :  clos
Fabric Health                           :  Black
 Fabric Status                          :  configure-success
 Fabric Level Physical Topology Health  :  Green
Fabric Device Health
+-------------+-------+--------------------+-------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+--------------------+-------------------+---------------+
| 10.20.246.6 | Leaf  | Black               | Black             | Black         |
```

```
| 10.20.246.5 | Leaf  | Black               | Black             | Black         |
| 10.20.246.1 | Spine | Black               | Black             | Black         |
| 10.20.246.2 | Spine | Black               | Black             | Black         |
| 10.20.246.4 | Leaf  | Green               | Green             | Green         |
| 10.20.246.3 | Leaf  | Green               | Green             | Green         |
+-------------+-------+---------------------+-------------------+---------------+

================================================================================================
================================================================
--- Time Elapsed: 49.241285ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
------------------------------------------------------------------------------------------------
----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.6 [Leaf]
  Device Health                          :  Black
  Configuration State Health             :  Black
    Dev State                            :  provisioned
    App State                            :  cfg refresh error
  Operational State Health               :  Black
    Cluster Health                       :  Black
      Operational State                  :  true
      Peer Operational State             :  false
      Peer Keepalive Operational State   :  false
    Physical Topology Device Health      :  Red
    Device physical topology errors
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+-----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+-----------------------------+---------------+
| 10.20.246.6    | Leaf             |                 |                       |
10.20.246.5         | Leaf                  |
|                            | missing-links |
+----------------+------------------+-----------------+-----------------------
+---------------------+-----------------------+----------------------
+-----------------------------+---------------+
    Underlay Topology Device Health      :  Black
    Device underlay topology errors
+------------------+----------------------+--------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+--------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+--------------------------+
| 10.20.246.6      | 10.20.246.5           | 172.31.254.55           | 10.20.20.15 |
```

```
65001             | 65001                  | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
+-----------------+-----------------------+-----------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.5 [Leaf]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Cluster Health                     :  Black
      Operational State                :  false
      Peer Operational State           :  false
      Peer Keepalive Operational State  :  false
    Physical Topology Device Health    :  Red
    Device physical topology errors
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
| 10.20.246.5    | Leaf             |                 |                      |
10.20.246.6         | Leaf                 |
|                           | missing-links |
| 10.20.246.5    | Leaf             |                 |                      |
10.20.246.1         | Spine                |
|                           | missing-links |
| 10.20.246.5    | Leaf             |                 |                      |
10.20.246.2         | Spine                |
|                           | missing-links |
+---------------+-----------------+----------------+----------------------
+--------------------+----------------------+---------------------
+----------------------------+---------------+
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+----------------------+-----------------------+-------------
+-------------------+-----------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+----------------------+-----------------------+-------------
+-------------------+-----------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.5      | 10.20.246.1           | 172.31.254.156          | 10.10.10.32 |
65001             | 64512                 | default-vrf | ipv4               |
unicast       |               | session_not_established |
| 10.20.246.5      | 10.20.246.1           | 172.31.254.156          | 10.10.10.32 |
65001             | 64512                 | default-vrf | l2vpn              |
evpn          |               | session_not_established |
| 10.20.246.5      | 10.20.246.2           | 172.31.254.156          | 10.10.10.34 |
65001             | 64512                 | default-vrf | ipv4               |
unicast       |               | session_not_established |
| 10.20.246.5      | 10.20.246.2           | 172.31.254.156          | 10.10.10.34 |
65001             | 64512                 | default-vrf | l2vpn              |
evpn          |               | session_not_established |
| 10.20.246.5      | 10.20.246.6           | 172.31.254.156          | 10.20.20.14 |
65001             | 65001                 | default-vrf | ipv4               |
```

```
unicast       |               | session_not_established |
+-----------------+-----------------------+-----------------------+-------------
+-------------------+-----------------------+-------------+--------------------
+---------------+---------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Physical Topology Device Health    :  Red
    Device physical topology errors
+----------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
| 10.20.246.1    | Spine           |                 |                      |
10.20.246.5         | Leaf                  |
|                          | missing-links |
+----------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
    Underlay Topology Device Health    :  Black
    Device underlay topology errors
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
| 10.20.246.1      | 10.20.246.5           | 172.31.254.144          | 10.10.10.33 |
64512             | 65001                 | default-vrf | ipv4              |
unicast       | CONN           | session_not_established |
| 10.20.246.1      | 10.20.246.5           | 172.31.254.144          | 10.10.10.33 |
64512             | 65001                 | default-vrf | l2vpn             |
evpn          | CONN           | session_not_established |
+------------------+-----------------------+-------------------------+-------------
+------------------+-----------------------+-------------+--------------------
+---------------+----------------+-------------------------+
 ------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Black
  Configuration State Health           :  Black
    Dev State                          :  provisioned
    App State                          :  cfg refresh error
  Operational State Health             :  Black
    Physical Topology Device Health    :  Red
    Device physical topology errors
+----------------+-----------------+----------------+----------------------
+---------------------+-----------------------+----------------------
+----------------------------+---------------+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE |     ERROR     |
+----------------+-----------------+----------------+----------------------
```

```
+---------------------+-----------------------+----------------------
+---------------------------+---------------+
| 10.20.246.2    | Spine            |                 |                       |
10.20.246.5         | Leaf                 |
|                            | missing-links |
+----------------+------------------+-----------------+----------------------
+---------------------+-----------------------+----------------------
+---------------------------+---------------+
    Underlay Topology Device Health     :  Black
    Device underlay topology errors
+------------------+-----------------------+-------------------------+-------------
+-------------------+------------------------+------------+-------------------
+---------------+----------------+------------------------+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN |     VRF     | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+------------------+-----------------------+-------------------------+-------------
+-------------------+------------------------+------------+-------------------
+---------------+----------------+------------------------+
| 10.20.246.2      | 10.20.246.5           | 172.31.254.205          | 10.10.10.35 |
64512             | 65001                  | default-vrf | ipv4               |
unicast       | CONN           | session_not_established |
| 10.20.246.2      | 10.20.246.5           | 172.31.254.205          | 10.10.10.35 |
64512             | 65001                  | default-vrf | l2vpn              |
evpn          | CONN           | session_not_established |
+------------------+-----------------------+-------------------------+-------------
+-------------------+------------------------+------------+-------------------
+---------------+----------------+------------------------+
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
------------------------------------------------------------------------------------------
---------------------------------------------------------------------

==========================================================================================
=====================================================================
--- Time Elapsed: 66.449895ms ---
```

After 2-3 mins, reloaded device comes up in a ready state.

```
(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Green
+-------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  | DEVICE STATE |  APP STATE  | CONFIG GEN
```

```
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+-------------+-------------
+------------------+----------------+--------+-------+
| 10.20.246.1 |     | NH-1      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.2 |     | NH-2      | 64512 | spine | provisioned  | cfg in-sync |
NA                | NA              | NA      | 1     |
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf  | provisioned  | cfg in-sync |
NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+-------------+-------------
+------------------+----------------+--------+-------+




(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
=====================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health  :  Green
Fabric Device Health
+-------------+-------+--------------------+------------------+---------------+
| IP ADDRESS  | ROLE  | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+-------+--------------------+------------------+---------------+
| 10.20.246.6 | Leaf  | Green              | Green            | Green         |
| 10.20.246.5 | Leaf  | Green              | Green            | Green         |
| 10.20.246.1 | Spine | Green              | Green            | Green         |
| 10.20.246.2 | Spine | Green              | Green            | Green         |
| 10.20.246.4 | Leaf  | Green              | Green            | Green         |
| 10.20.246.3 | Leaf  | Green              | Green            | Green         |
+-------------+-------+--------------------+------------------+---------------+

==========================================================================================
=====================================================================
--- Time Elapsed: 37.195902ms ---



(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
==========================================================================================
=====================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Green
 Fabric Status                            :  configure-success
 Fabric Level Physical Topology Health   :  Green
------------------------------------------------------------------------------------------
---------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.6 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
```

```
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.5 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg refreshed
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.1 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg refreshed
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.2 [Spine]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg refreshed
  Operational State Health             :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.4 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
 --------------------------------------------------------
 Device IP [Role]                      :  10.20.246.3 [Leaf]
  Device Health                        :  Green
  Configuration State Health           :  Green
    Dev State                          :  provisioned
    App State                          :  cfg in-sync
  Operational State Health             :  Green
    Cluster Health                     :  Green
    Physical Topology Device Health    :  Green
    Underlay Topology Device Health    :  Green
------------------------------------------------------------------------------------------
---------------------------------------------------------------------

==========================================================================================
=====================================================================
--- Time Elapsed: 34.651481ms ---
```

Verify that Spine Devices Have Been Deleted from the Fabric

The following sample output verifies the deletion of spine devices from the Fabric:

```
(efa:user)user@dev-server:~$ efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: configure-success, Fabric Health: Red
+-------------+-----+-----------+-------+------+--------------+-------------
+------------------+----------------+--------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE | DEVICE STATE |  APP STATE  | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+------+--------------+-------------
+------------------+----------------+--------+-------+
| 10.20.246.3 |     | NH-Leaf1  | 65002 | leaf | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
| 10.20.246.4 |     | NH-Leaf2  | 65002 | leaf | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
| 10.20.246.5 |     | NHF-Leaf1 | 65001 | leaf | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
| 10.20.246.6 |     | NHF-Leaf2 | 65001 | leaf | provisioned  | cfg in-sync |
NA               | NA              | 2       | 1     |
+-------------+-----+-----------+-------+------+--------------+-------------
+------------------+----------------+--------+-------+

CONFIG GEN REASON:
LA/LD - Link Add/Delete, IA/ID/IU - Interface Add/Delete/Update, PLC/PLD/PLU -
IPPrefixList Create/Delete/Update
MD/MU - MCT Delete/Update, OD/OU - Overlay Gateway Delete/Update, EU/ED - Evpn Delete/
Update, PC/PD/PU - RouterPim Create/Delete/Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn Update, SYS - System
Properties Update
MD5 - BGP MD5 Password, BGPU - Router BGP Update, BGPLL - BGP Listen Limit, POU - Port
Channel Update, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties, INTIP - Interface IP,
BGP - Router BGP
C/D/U - Create/Delete/Update, PA/PD - Port Add/Port Delete

For App or Device Error/Failure reason, run "efa fabric error show" for details
For config refresh reason, run "efa fabric debug config-gen-reason" for details

--- Time Elapsed: 54.992785ms ---


(efa:user)user@dev-server:~$ efa fabric health show --name fab3
==========================================================================================
=====================================================================
Fabric Name                           :  fab3
Fabric Type                           :  clos
Fabric Health                         :  Red
 Fabric Status                        :  configure-success
 Fabric Level Physical Topology Health  :  Red
Fabric Device Health
+-------------+------+---------------------+-------------------+---------------+
| IP ADDRESS  | ROLE | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-------------+------+---------------------+-------------------+---------------+
| 10.20.246.6 | Leaf | Green               | Green             | Green         |
| 10.20.246.5 | Leaf | Green               | Green             | Green         |
| 10.20.246.4 | Leaf | Green               | Green             | Green         |
| 10.20.246.3 | Leaf | Green               | Green             | Green         |
+-------------+------+---------------------+-------------------+---------------+
```

```
================================================================================================
================================================================
--- Time Elapsed: 34.489515ms ---
(efa:user)user@dev-server:~$ efa fabric health show --name fab3 --detail
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Red
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Red
------------------------------------------------------------------------------------------------
----------------------------------------------------------------
Fabric level topology errors     :
+---------------------+----------------+---------------+
| MISSING SUPERSPINES | MISSING SPINES | MISSING LEAFS |
+---------------------+----------------+---------------+
| true                | true           | false         |
+---------------------+----------------+---------------+
------------------------------------------------------------------------------------------------
----------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.20.246.6 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.5 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.4 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
 ------------------------------------------------------
 Device IP [Role]                        :  10.20.246.3 [Leaf]
  Device Health                          :  Green
  Configuration State Health             :  Green
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Green
    Cluster Health                       :  Green
    Physical Topology Device Health      :  Green
    Underlay Topology Device Health      :  Green
------------------------------------------------------------------------------------------------
----------------------------------------------------------------
```

```
================================================================================================
================================================================
--- Time Elapsed: 32.372653ms ---
```

Remove all Devices from Fabric

the following sample output removes all devices from fabric:

```
(efa:root)root@admin01:~# efa fabric show --name fab3

Fabric Name: fab3, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: created, Fabric Health: Green
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE | DEVICE STATE | APP STATE | CONFIG GEN
REASON | PENDING CONFIGS | VTLB ID | LB ID |
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+
--- Time Elapsed: 203.154048ms ---
(efa:root)root@admin01:~# efa fabric health show --name fab3
================================================================================================
================================================================
Fabric Name                              :  fab3
Fabric Type                              :  clos
Fabric Health                            :  Green
 Fabric Status                           :  created
 Fabric Level Physical Topology Health   :  Green

================================================================================================
================================================================
--- Time Elapsed: 170.378237ms ---
(efa:root)root@admin01:~# efa fabric health show --name fab3 --detail
================================================================================================
================================================================
Fabric Name                               :  fab3
Fabric Type                               :  clos
Fabric Health                             :  Green
 Fabric Status                            :  created
 Fabric Level Physical Topology Health    :  Green
------------------------------------------------------------------------------------------------
----------------------------------------------------------------


================================================================================================
================================================================
--- Time Elapsed: 139.848544ms ---
```

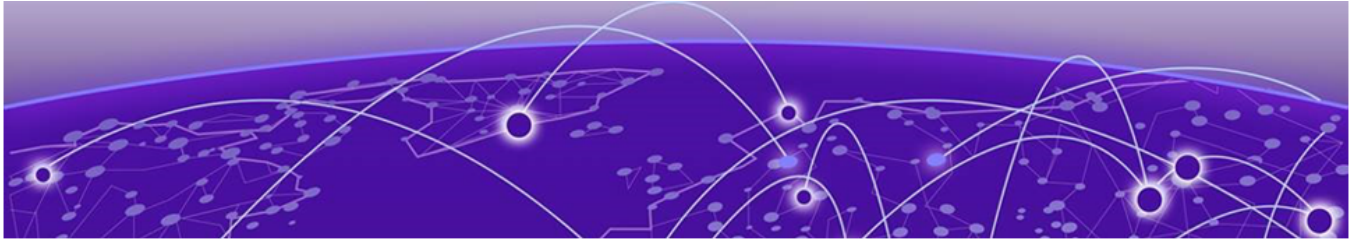

Known Limitations

Note the following caveats for ExtremeCloud Orchestrator.

Known Limitations in Fabric Skill

Follow these caveats and limitations when using Fabric Skill.

VRF delete from EPG and re-adding VRF to EPG fails intermittently

Symptom	Condition	Workaround
Endpoint group (EPG) update **vrf-add** operation fails with the reason as VRF to be added has conflicting VRF on the switch.	Run EPG update **vrf-add**, **vrf-delete**, and **vrf-add** operation CLI in quick succession: 1. Update EPG for operation **vrf-add**. 2. Update EPG for operation **vrf-delete**. 3. Update the same EPG again with operation **vrf-add** for the same VRF which was deleted in step 2.	Wait of 30 seconds between the EPG update **vrf-add** and vrf-delete operations on the same EPG.

REST operations are not retried (as applicable) during the service boot

Symptom	Condition	Workaround
REST operations are not retried (as applicable) during the service boot up.	After publishing the necessary events on the message bus, the status for the REST operations are nit set automatically.	Manually set the status for all REST operations.

RBAC: XCO shows "export EFA_TOKEN" command suggestion when a tenant user logs in

Symptom	Condition	Workaround
XCO shows a **export EFA_TOKEN** message after a tenant user with RBAC logs in to the system.	When a user is created with the default login shell as sh.	XCO supports only bash shell for login or any other CLI commands. Ensure that bash is specified as the default login shell for all XCO user accounts.

EFA Token command message

Please type this in your shell:

export EFA_TOKEN=eyJhbGciOiJSUzI1NiIsImtpZCI6IjEuMCIsInR5cCI6IkpXVCJ9.eyJjb21tb25fbmFtZSI6IkVGQSBUb2tlbiBTZXJ2aWNlIiwidWFzIjpbeyJ0YXJnZXQiOiJFRkEiLCJyb2xlIjoiVlIyLVRudEFkbWluIn1dLCJvcmciOiJFeHRyZW1lIE5ldHdvcmtzIiwidmVyIjoiMS4wIiwiaWQiOiIiLCJleHAiOjE2NDUyNDcxNDIsImp0aSI6IjZjMjA4ZDUxLTkwNzgtMTFlYy1iZjk5LWNhNzk1MDY1YzIwNyIsImlhdCI6MTY0NTE2MDc0MiwiaXNzIjoiRUZBIFRva2VuIFNlcnZpY2UiLCJuYmYiOjE2NDUxNjA3NDIsInN1YiI6InVzZXIyIn0.b7m5PINijeEdNSqnTeE2ZhUrqKLKQAu079vXyBIdgHbXKt9ULfa03vMU1jfBO1qFb1-x0oHmsAQ0pSsF5JLeMaMzMflLf78ktZO8U5IePq72vM5en35IR-DNLyoGIZBeFeG6ZbBMoETzz5vf9OuefgQID3YdjcALr7yllCgDmLVFlgson77yCBpkTK15xm1GRbtL7JKXZzShBE7E3kdW7N71MdM85Gc3r4l-c8sfz7eo06gKrfTq9wXCv4_LVzR6-KRSg6NyLq363WEpcK1A2Hs0Wo3T9TpquYHNaCWA5I1QTsG-RHFdg4kxZP2fQpUp6Bgy1s6k59PVPn4-M-a8lA— Time Elapsed: 4.619465187s ——

XCO CLI or REST request with scale config takes longer than 15 minutes fails

Symptom	Condition	Workaround
Tenant2 delete is successful whereas deleting Tenant1 took more than 15 minutes and failed with the following message: Error : service is not available or internal server error has occurred, please try again later Tenant service was running. Tenant1 was not available after the error.	When you try to delete tenants in a single rack small data center deployment configured with scale tenant config	Any CLI or REST tenant operations, and any fabric operations taking more than 15 minutes, will timeout at the client side. The operation completes in the background. Run the **efa tenant show** command to view the actual state of the operation.