ExtremeCloud™ Orchestrator v4.1.0 CLI Administration Guide

CLI Configuration, Management, and Troubleshooting

9041107-00 Rev AA
September 2026

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: https://www.extremenetworks.com/about-extreme-networks/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/

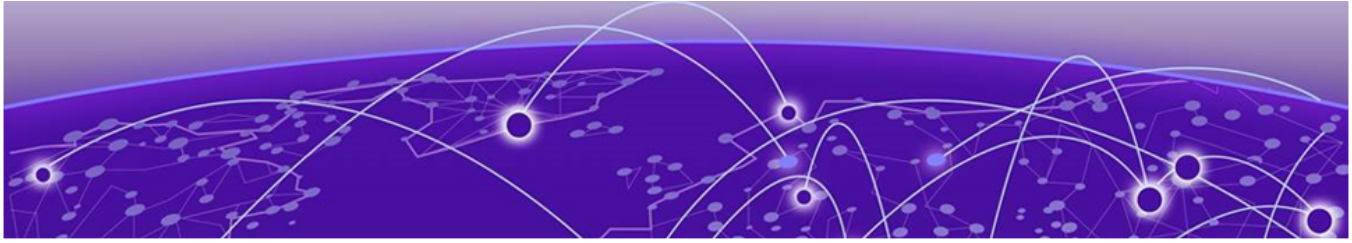

Table of Contents

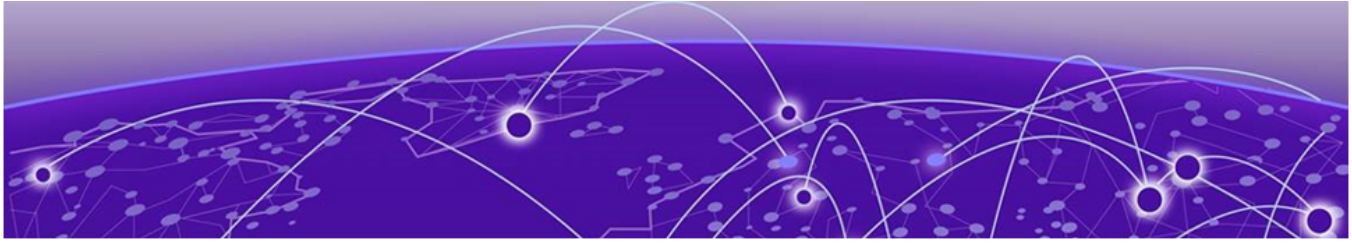

Abstract

This CLI administration guide for ExtremeCloud™ Orchestrator version 4.1.0 provides detailed technical instructions for configuring, managing, and troubleshooting the XCO platform using command-line interface tools. Designed for IT administrators, the guide outlines procedures for system verification, certificate lifecycle management, tenant and fabric provisioning, and high-availability deployments. It introduces XCO's microservices-based architecture and integration with Extreme Fabric Automation (EFA), supporting SLX and Extreme OS ONE devices. Key features include automated backup and restore, advanced logging and audit trail configuration, and support for REST APIs and SSH key-based authentication. The document also details the configuration of Clos and non-Clos data center fabrics, tenant VRFs, endpoint groups, BGP peers, and traffic mirroring in multi-tenant environments. Version 4.1.0 introduces updates such as next-hop route-map support, bulk API enhancements for tenant EPGs, and improved supportsave diagnostics. The guide is technically advanced and suitable for experienced network engineers managing scalable, secure, and automated data center networks.

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 Extreme Networks switches, the product is referred to as *the switch*.

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 articles, white papers, and case studies

Open Source Declarations

Some software files have been licensed under certain open source licenses. Information is available on the Open Source Declaration page.

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit the Extreme Networks Training page.

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 2800. 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 User Enablement 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, email us at Product-Documentation@extremenetworks.com.

Provide as much detail as possible including the publication title, topic heading, and page number (if applicable), along with your comments and suggestions for improvement.

What's New in this Document

The following table describes information added to this guide for the XCO 4.1.0 software release.

Table 4: Summary of changes

Feature	Description	Link
XCO MPLS L3 VPN Inter AS Support	New topic "Configure MPLS L3 VPN Inter AS" describes the procedure to configure Inter AS between SLX 9740 and PE.	Configure MPLS L3 VPN Inter AS on page 568
IP Fabric Underlay BFD Session Transition	New topic "BFD Session Transition on IP Fabric Underlay" and its sub-topics describe the architecture and step-by-step provisioning guidance for executing the BFD session migration.	BFD Session Transition on IP Fabric Underlay on page 247 • Fabric Setting Update — BFD Session Type Flag on page 255 • XCO Automation, Versioning, and Upgrade on page 257 • User-Initiated Transition to Software BFD on page 259 • Provision Fabric with BFD Configuration on page 260 • Compliance, Security, and Miscellaneous on page 272

Table 4: Summary of changes (continued)

Feature	Description	Link
Fabric Health Enhancements	Updated the following topics: • "Health Management" topic with brief details on new enhancements • "Health Status Levels" table in the "Bubbling of Health Status" topic • "Fabric Health" topic with the fabric status and its aggregation rules. • "Health APIs" table with more REST APIs • "Fabric Health Calculation" topic with a new "Triggering a Health Recalculation" table • Updated the table with new commands in the "Health Command" topic	Health Management on page 1019 • Bubbling of Health Status on page 1020 • Fabric Health on page 1024 • Health APIs on page 1023 • Fabric Health Calculation on page 1025 • Health Commands on page 1022
	Added the following topics: • Service Health • Fabric Health Failure Scenarios — 5-Clos Topology (Two PODs) • Fabric Health Failure Scenarios — 3-Clos Topology (Two MCT Pairs)	• Service Health on page 1033 • Fabric Health Failure Scenarios — 3-Clos Topology (Two MCT Pairs) on page 1037 • Fabric Health Failure Scenarios — 5-Clos Topology (Two PODs) on page 1042
Certificate Management Improvements — CMPv2 and Revocation	Added a new topic "Certificate Management Using CMPv2" and sub-topics that describe configuration and management of certificates in XCO using CMPv2.	Certificate Management Using CMPv2 on page 62 • EJBCA CMPv2 Pre-Configuration Requirements on page 64 • Configure CMPv2 on page 65 • Auto Renewal of Certificate on page 68 • Switch Certificate Management Mechanism on page 69 • Compatibility and Upgrade Considerations on page 70

Table 4: Summary of changes (continued)

Feature	Description	Link
Switch Security Settings — Minimum TLS Version	Updated the topic "Minimum TLS Version" with supported platforms, TLS versions, services for override (OS ONE), configuration behaviour and validation. New topic "How TLS Version Configuration Works" describes the working behavior of TLS configuration on SLX-OS and OS ONE devices. New topic "Configure Minimum TLS Version on Device" describes the procedure to configure minimum TLS version on SLX-OS and OS ONE devices.	• Minimum TLS Version on page 85 • How TLS Version Configuration Works on page 89 • Configure Minimum TLS Version on Device on page 90
Firmware Upgrade Auto-Recovery	New topics "Firmware Upgrade Auto-Recovery", "Firmware Upgrade Auto-Recovery Flow", and "Firmware Upgrade Failure Scenarios" describe the conceptual information, failure scenarios, and auto-recovery flow.	Firmware Upgrade Auto Recovery on page 767 • Firmware Upgrade Failure Scenarios on page 769 • Firmware Upgrade Auto Recovery Flow on page 770
XCO Drop Precedence Support for Extreme OS ONE	New topics "Drop Precedence Support for Extreme OS ONE" and "Configure QoS Maps with Drop Precedence" describe conceptual and procedural information to configure drop precedence for Extreme OS ONE.	• Drop Precedence Support for Extreme OS ONE on page 701 • Configure QoS Maps with Drop Precedence on Extreme OS ONE on page 702
Prefix list support on Extreme OS ONE	New topics "Prefix List for Extreme OS ONE Devices" and "Configure IP Prefix List on Extreme OS ONE Devices" describe the prefix list configuration details on Extreme OS ONE devices.	• Prefix List for Extreme OS ONE Devices on page 642 • Configure IP Prefix List on Extreme OS ONE Devices on page 644

Table 4: Summary of changes (continued)

Feature	Description	Link
Route map support on Extreme OS ONE	New topics "Route Map for Extreme OS ONE Devices", "Configure Route Map on Extreme OS ONE Devices", "Configure Route Map Match Condition on Extreme OS ONE Devices", and "Configure Route Map Set Action on Extreme OS ONE Devices" describe the route map configuration details on Extreme OS ONE devices.	• Route Map for Extreme OS ONE Devices on page 653 • Configure Route Map on Extreme OS ONE Devices on page 657 • Configure Route Map Match Condition on Extreme OS ONE Devices on page 658 • Configure Route Map Set Action on Extreme OS ONE Devices on page 659
Firmware Upgrade on OS ONE Devices	The following topics are updated: • Firmware Upgrade on OS ONE Devices • CLI Commands for Firmware Upgrade on OS ONE Devices • Customized Upgrade Flows The following new topics are added: • Upgrade Firmware on OS ONE Devices: Describes procedure for upgrading firmware on OS ONE devices • Monitor Firmware Upgrades: Describes the procedure to monitor firmware upgrade status and maintain the historical data. • Roll Back Firmware Upgrade: Describes the procedure for rollback the firmware upgrade. • Maintenance Mode and DRC: Describes the maintenance mode and DRC for firmware upgrades.	• Firmware Upgrade on OS ONE Devices on page 742 • CLI Commands for Firmware Upgrade on OS ONE Devices on page 744 • Customized Upgrade Flows on page 745 • Upgrade Firmware on OS ONE Devices on page 746 • Monitor Firmware Upgrades on page 748 • Roll Back Firmware Upgrade • Maintenance Mode and DRC on page 748

Table 4: Summary of changes (continued)

Feature	Description	Link
BGP peer group and route map configuration on OS ONE devices	New topic "BGP Peer Groups with Route Map Policies on OS ONE Devices" and sub-topics describes BGP peer group, BGP peer, and route map configuration on OS ONE devices.	BGP Peer Groups with Route Map Policies on OS ONE Devices on page 540 • Configure BGP Peer Groups on OS ONE Devices on page 542 • Configure BGP Peers on OS ONE Devices on page 544 • Configure Route Maps on OS ONE Devices on page 544 • Common Configuration Patterns: Configure Customer BGP Session with Import/Export Control on page 545
Fabric Health in XCO	The following new topics are added: • Fabric Health Composition: Describes device-level and fabric-level health compositions. • Data Sources and Authority Model: • Health Data Collection and Reconciliation: Describes the use of hybrid model to ensure timely updates and eventual consistency. • Known Behavioral Considerations and Enhancements: Describes the known issues and improvements. The following topic is updated: • Unified Health and Fault Management Overview: Added conceptual information about fabric health.	• Fabric Health Composition on page 917 • Data Sources and Authority Model on page 917 • Health Data Collection and Reconciliation on page 918 • Known Behavioral Considerations and Enhancements on page 919
Trace-level Logging	Updated the topic "Configure Logging" with the CLI to enable trace-level logging.	Configure Logging on page 120

Table 4: Summary of changes (continued)

Feature	Description	Link
Fabric health cron interval	New topics "Health Check Interval" and "Configure Fabric Health Cron Interval" describe the health check interval and the procedure to view and update the fabric health cron job interval using CLI and REST API.	Health Check Interval on page 1033 • Configure Fabric Health Cron Interval on page 1034
SNMP Trap and Delivery	A now topic "SNMP Trap & Inform Delivery" describes SNMP traps and informs configuration using XCO proxy path, direct path, and hybrid models.	SNMP Trap & Inform Delivery on page 877
Health Color Precedence Order	Updated the topic "Bubbling of Health Status" with the details of health color precedence order	Bubbling of Health Status on page 1020
Interface Configuration Behavior Changes	New topic "Interface Configuration Behavior Changes" describes the interface configuration tracking improvements introduced in XCO v4.1.0.	Interface Configuration Behavior Changes on page 786
MD5 Password Encoding	New topic "MD5 Password Encoding in Fabric Settings" describes the encoding behavior of the BGP MD5 password before storing it in the XCO database.	MD5 Password Encoding in Fabric Settings on page 216
Recover Rogue Configuration	The new topics "Recover XCO and End-Device Configuration" describes the procedure to recover XCO and end-device configuration.	Recover XCO and End-Device Configuration on page 100

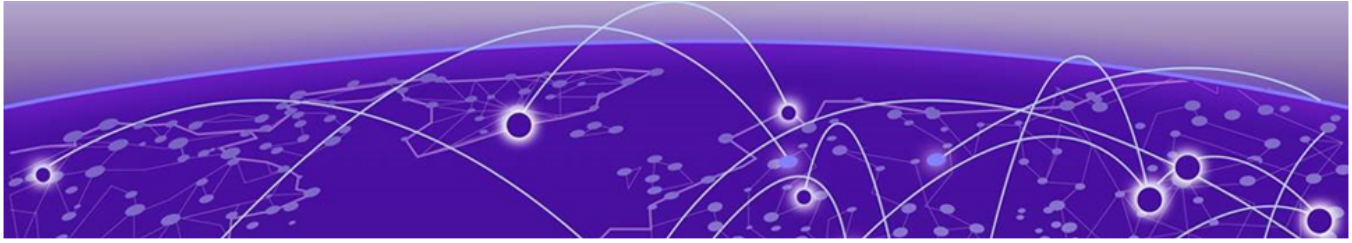

Introduction to ExtremeCloud Orchestrator

ExtremeCloud™ Orchestrator (XCO) is a single layer orchestration application that provides a unified and holistic GUI and APIs for fabric-wide life cycle management with highly scalable and flexible deployment model for Extreme solutions.

XCO provides common infrastructure and consistent installation and upgrade strategies for SLX, and 8000 series 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 policy management of SLX devices.

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). 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 26.

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

XCO Architecture

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

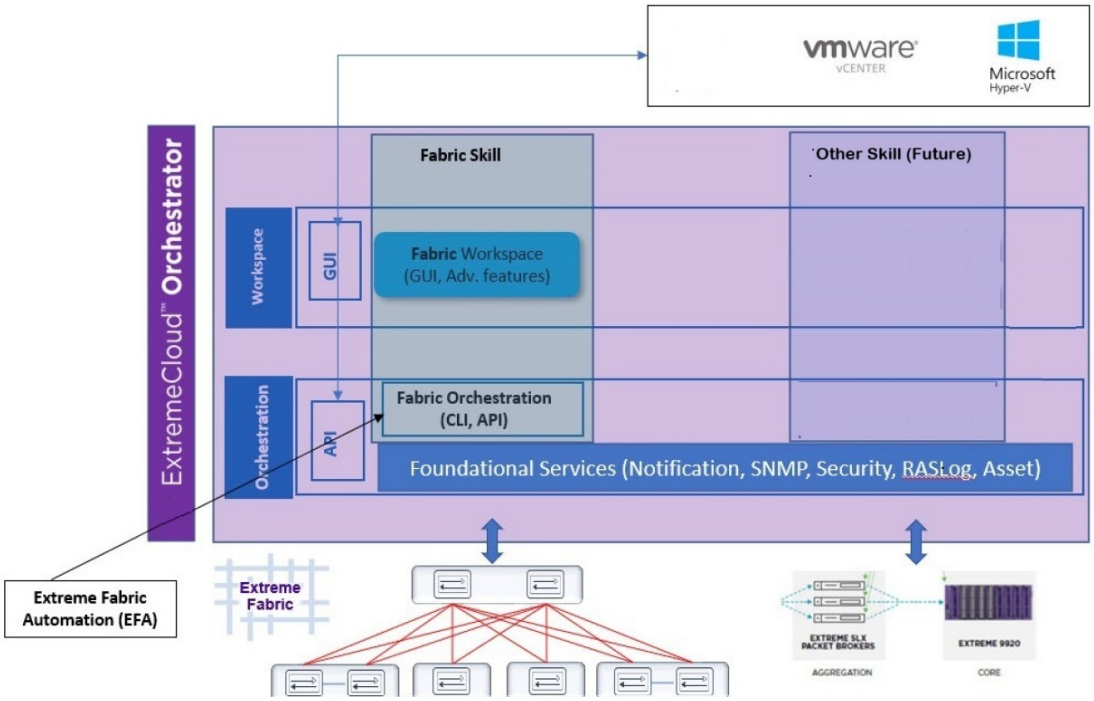

Figure 1: Overview of ExtremeCloud Orchestrator

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 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 28.

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, 4.1.0* and REST API Documentation for XCO on page 32.

Deployment

For more information about deployment scenarios, see the *ExtremeCloud Orchestrator Deployment Guide, 4.1.0*.

XCO on TPVM

TPVM (Third-Party Virtual Machine) is a guest VM that resides on Extreme SLX devices. You can run XCO from the Extreme 8520, Extreme 8720, Extreme 8820, 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.

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.

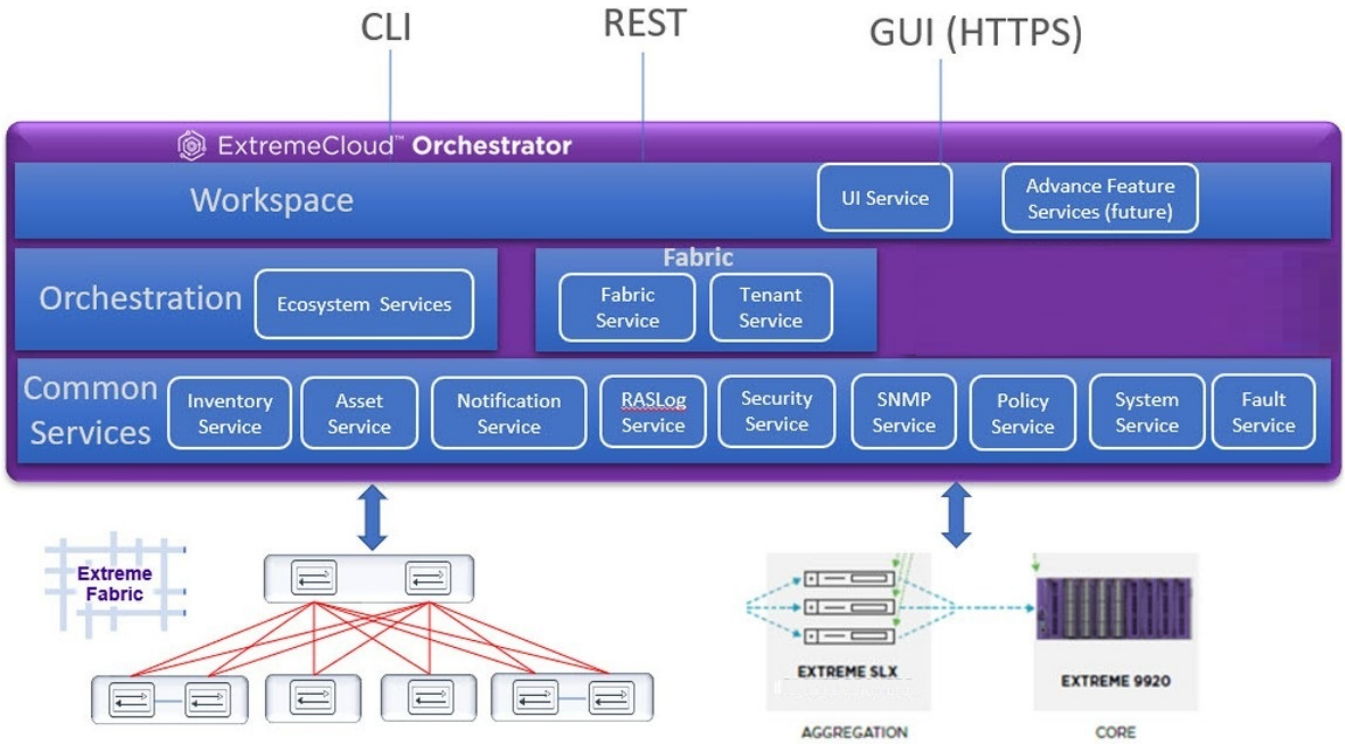

Figure 2: Microservices in the XCO architecture

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, 4.1.0*.

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.

You can use LDAP with XCO for user authentication and authorization. Based on the server configuration, XCO provides various options to configure LDAP.

- TLS Service: The service enables encrypted communication from XCO to LDAP server.
- 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.

For more information, see XCO User Authentication and Authorization in the *ExtremeCloud Orchestrator Security Configuration Guide, 4.1.0*

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, 4.1.0*.

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.

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 *ExtremeCloud Orchestrator VMware vCenter Integration Guide, 4.1.0*. 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 *ExtremeCloud Orchestrator Hyper-V Integration Guide, 4.1.0*. Integration support includes the following:

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

XCO OS ONE Integration and Platform Expansion

XCO 4.0.2 and later releases support for XCO OS ONE devices and adds compatibility with the Extreme 8730, 8720, and 8820 platforms.

Note

- The features related to XCO OS ONE and Extreme 8730 support in XCO 4.1.0 are released as Control Released Features.
- For comprehensive feature descriptions, CLI command references, administrative procedures, and API documentation, contact Extreme Networks.

Extreme OS ONE Overview

Extreme OS ONE is a cloud-native, microservices-based network operating system designed for IP fabrics and data center environments. Key features include:

- **Modular Architecture**: Separates Base OS and Application OS for flexible deployment
- **API-First Design**: Supports gNMI/gNOI with OpenConfig YANG models and CLI.
- **Security-First Principles**: Includes AAA integration, X.509 certificate management, and secure access protocols
- **Lifecycle Management**: Enables firmware install, rollback, commit, diagnostics, and ZTP
- **Monitoring & Logging**: Offers extensive telemetry, syslog support, and threshold-based alerts
- **Integrated Appliance Hosting**: Supports third-party VM deployment via KVM

Supported Platforms

Extreme 8730 (running Extreme OS ONE)

Extreme 8730 Platform

The Extreme 8730-32D is a high-performance switching and routing platform powered by Extreme OS ONE. Its capabilities include:

- **Port Operations**: 400G support with breakout, RS-FEC, and link diagnostics
- **Layer 2 Features**: MLAG, LAG, bridge domains with VLAN tagging and MAC learning
- **Layer 3 Routing**: Full IPv4/IPv6 stack, BGP (iBGP/eBGP), ECMP, and VXLAN/EVPN overlays
- **Resiliency**: BFD, resource threshold monitoring, and thermal/fan policy management

Key XCO Integration Capabilities

XCO offers the following key integration capabilities:

- Device registration and inventory management.
- Secure settings enforcement via gNMI.
- Interface and system configuration.
- Fabric Integration for IP Fabric deployments.
- Drift and reconcile operations

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
- RASlog service
- RBAC service

- SNMP service
- System service
- Tenant service
- vCenter service
- Policy service

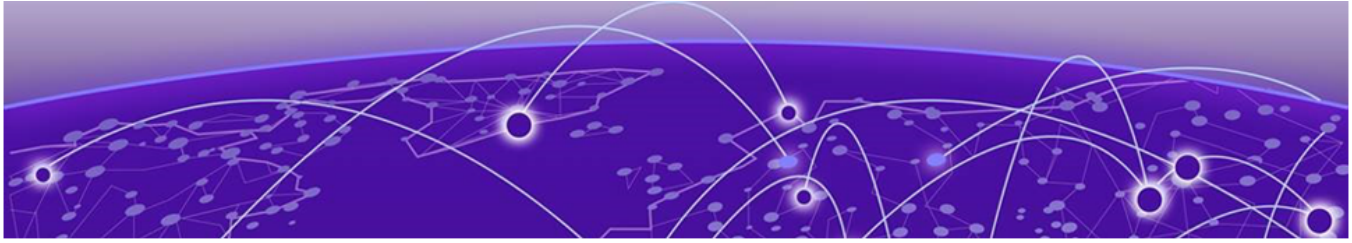

XCO System Management

Learn about configuring system-level settings such as supportsave, backup and restore, management routes, logging, and certificates.

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-OS command prompt.

      Ensure that the status of all the nodes are up.

```
device# show efa status
=================================================
               EFA version details
=================================================
Version : 3.4.0
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.

Login Banner

In today's digital landscape, securing access to sensitive systems and data is a top priority for any organizations. A key element of this is the use of security banners (banner text) or login banners on login screens at a hierarchy level. The banner text appears during SSH CLI logins, console logins, and GUI logins. These banners are in compliance with government and company policies and inform end-users or admins of the terms of use, monitoring policies, and legal implications of unauthorized access.

Security banners play a crucial role in protecting systems, ensuring legal compliance, and demonstrating an organization's commitment to cyber security. By clearly communicating expectations, they help deter unauthorized activities and enhance user accountability.

When a banner is configured, users will see the banner displayed on the XCO system (including TPVM and OVA) login page after they enter their credentials and when they click the **Login** button. An **Agree** and **Cancel** button is shown beneath the banner. Users then need to click the **Agree** button to complete the login. If they click **Cancel**, they are returned to the login page.

XCO allows you to set a login banner for the XCO system, applicable to SSH and console logins, as well as XCO application CLI and GUI logins. During the XCO deployment, a default banner is configured. This banner will appear when the users re-login or open a new session. The login

banner is part of the SSH configuration within the system and is stored in a designated file. The following table describes the file locations based on the deployment type:

Deployment Type	Location
Server/OVA	/opt/efadata/misc/security_banner
TPVM	/apps/efadata/misc/security_banner

In a multi-node setup, both nodes will display the same banner message. This is achieved by storing the banner file in a location shared among the nodes using GlusterFS. When the user configures the banner on one XCO node, it is automatically applied to both nodes.

XCO triggers an SSHD restart in the following scenarios:

- XCO upgrade
- Backup and restore of XCO
- TPVM upgrades

Note
XCO overwrites any existing login banner text, If a backup does not contain a banner, and the restored version will also not include it.

SSHD Configuration

The path to the banner file is configured in the SSHD config file and requires a restart if the path changes. This restart occurs during the XCO deployment.

The following line is added to the **banner conf** file in the `sshd_config.d` directory: `Banner /apps/efadata/misc/security_banner`

When the security banner message is updated, there will be no changes needed in the `sshd_config.d/banner.conf` file since the file path remains the same.

The following tables outlines the scenarios in which XCO performs the restart of SSHD service:

Scenarios	Is SSHD restarted?
XCO fresh deployment (3.8.0 and later)	Yes
XCO fresh deployment (3.8.0 and later)	Yes (after removal of banner)
XCO upgrade (prior versions to 3.8.0 and later)	Yes
XCO upgrade (from 3.8.0 to later)	Yes
XCO upgrade (from 3.8.0 to later)	No
TPVM upgrade yes (only on the standby)	Yes

Console Login Configuration

To configure a security banner for console logins, ensure that the message is added to the `/etc/issue` file. In addition to updating the security_banner file used for SSH and XCO logins, the `/etc/issue` file will also be updated on both nodes to ensure consistency.

Note
If XCO is un-deployed, the /etc/issue file will revert to its default console login message, rather than any previous custom message.

Logging

The logs from CLI and Rest executions are included in the monitor logs, and are available at the following locations:

- Server : /var/log/efa/monitor
- TPVM : /apps/efa_logs/monitor

Banner Guidelines

When creating a banner, ensure it clearly communicates the following:

1. Unauthorized access prohibition: Explicitly state that unauthorized access is prohibited.
2. Activity monitoring and logging: Inform users that their activities will be monitored and logged.
3. Terms of use consent: Indicate that users consent to the terms of use by accessing the system.
4. Message length: Banner messages must not exceed 500 characters in length.

Configuration Best Practices

1. Clear and concise language: Avoid technical jargon or vague wording.
2. Accuracy and transparency: Ensure the banner does not contain misleading or incorrect information.
3. Sensitive information protection: Refrain from including sensitive or classified details in the banner.
4. Length and concision: Limit the banner to 20-30 lines to avoid cumbersome messages.
5. Relevant content: Focus on essential information, such as system status, important notices, or security warnings.
6. Performance optimization: Keep scripts generating the security banner lightweight to prevent login process delays.

Configure System-Wide Banner

You can display a customizable and system-wide message using banner text.

About This Task

Follow this procedure to configure a system-wide banner text using CLI and API.

Note

Only users with System Admin and Security Admin privileges can configure the banner.

Procedure

1. To configure a banner text, run the following command:

```
$ efa system security-banner set --banner-message
'WARNING: This system is for authorized use only. Activities on this system are
monitored and recorded. Unauthorized access or use is prohibited and may result in
disciplinary action and/or legal prosecution.'

Setting up the security banner is successful
```

   Alternatively, use the following API:

```
POST: /v1/system/security-banner
REQUEST BODY SCHEMA: application/json
{
"message": "Setting up the security banner is successful."
}
```

2. To display a banner text, run the following command:

```
$ efa system-security banner show

WARNING: This system is for authorized use only. Activities on this system are
monitored and recorded. Unauthorized access or use is prohibited and may result in
disciplinary action and/or legal prosecution.
```

   Alternatively, use the following API:

```
GET:
/v1/system/security-banner

RESPONSE SCHEMA: application/json

   {

"message": "WARNING: This system is for authorized use only. Activities on this system
are monitored and recorded. Unauthorized access or use is prohibited and may result in
disciplinary action and/or legal prosecution."

}
```

3. To reset a banner text, run the following command:

```
$ efa system security-banner reset

security banner reset is successful
```

   Alternatively, use the following API:

```
POST:
v1/system/security-banner/reset
```

```
RESPONSE SCHEMA: application/json

{

"message": Security banner reset is successful

}
```

4. To delete a banner text, run the following command:

```
$ efa system security-banner unset

Security banner has been deleted successfully
```

Alternatively, use the following API:

```
DELETE:
v1/system/security-banner


RESPONSE SCHEMA: application/json

{

HTTPS status code : 204 No content

}
```

Example

- TPVM or Server Login

  Following XCO installation, a security banner message is displayed when users attempt to access the TPVM/Server via SSH. Conversely, if XCO is not installed on the system, no banner message will appear.

```
user@test:~$ ssh extreme@10.20.54.91
WARNING: This system is for authorized use only. Activities on this system are
monitored and recorded. Unauthorized access or use is prohibited and may result in
disciplinary action and/or legal prosecution.
extreme@10.20.54.91's password:
```

- XCO Login

  When XCO is installed, a security banner message will be displayed when a user attempts to log in to the TPVM/Server using SSH. A banner message will not be shown if XCO is not installed on the system.

```
extreme@tpvm:~$ efa login
WARNING: This system is for authorized use only. Activities on this system are
monitored and recorded. Unauthorized access or use is prohibited and may result in
disciplinary action and/or legal prosecution.
Password:
```

- Console Login

  Upon console login to the server, a security banner is displayed with the following message:

```
WARNING: This system is for authorized use only. Activities on this system are
monitored and recorded. Unauthorized access or use is prohibited and may result in
disciplinary action and/or legal prosecution.
$ efa login:
```

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 and Extreme OS ONEdevice 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.

When registering a Extreme OS ONE device in XCO, the following certificates are installed on the device:

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

Along with the certificate installation, the following configuration changes are done on the registered SLX 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`.

Along with the certificate installation, the following configuration change is done on the registered OS ONE device:

1. Configure the grpc-server
2. Assign the certificate ID

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=<username> --password=<password>
+----+------------+-----------+-------+--------------+----------+---------+--------+
| ID | IP Address | Host Name | Model | Chassis Name | Firmware | Status  | Reason |
+----+------------+-----------+-------+--------------+----------+---------+--------+
```

```
| 1  | 10.x.x.x   | SLX       | 3012  | Extreme 8720-32C| 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 922.

Performing a manual device update or DRC will also update the SLX syslog certificate.

- Manual Device Update

  ```
  efa inventory device update (--ip [ip_address] | --fabric [fabric_name])
  ```

- Manual Drift Reconcile

```
efa inventory drift-reconcile execute --ip [ip_address] --reconcile
```

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 922.

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 46.

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

HTTPS Certificates

When you register a SLX 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 SLX devices and GRPC server on the OS
ONE devices. The 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 922.

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 SLX devices and GRPC server on the OS
ONE devices. The 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 46.

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

GRPC Certificates

When you register an OS ONE device in XCO, a new certificate is generated for the GRPC server of the OS ONE device. The certificate is generated with the default CA that XCO contains. You can use the CLI command to install third-party certificates on a single OS ONE device at once.

```
$ efa certificate device install --ip 10.x.x.x --grpc-certificate device1.pem --grpc-key
device1.key

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

Do you want to proceed [y/n]
+---------------+---------+
| IP Address    | Status  |
+---------------+---------+
| 10.x.x.x      | Success |
+---------------+---------+
```

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
      --grpc-certificate string    Local path to the gRPC certificate pem file
      --grpc-key string            Local path to the gRPC 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

All of the XCO components produce and use different certificates.

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.
8. **Galera Certificate**: XCO uses Mariadb database with galera cluster for replication. This certificate enables SSL for the replication across nodes. This is only applicable for multi-node installation of XCO.

The following tables provide information about XCO certificates.

For Alerts related to Alarms or Notifications, see Fault Management - Alerts on page 922.

SSL/TLS Certificate of XCO

Location in TPVM deployment	/apps/efadata/certs/own/tls.crt
Location in server deployment	/opt/efadata/certs/own/tls.crt
Description	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)
Default Validity Period	Expires in 3 years from installation. Reset after every subinterface creation/upgrade
Impact on the system	If the certificate expires, then the server communication with SSL verification enabled will fail. Disables syslog messages from the devices
Renewal Procedure	Use the **`efa certificate server renew`** command as described in the XCO Server Certificate on page 54.
Alarm/Notification	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

Location in TPVM deployment	/apps/rancher/k3s/server/tls/server-ca.crt
Location in server deployment	/var/lib/rancher/k3s/server/tls/server-ca.crt
Description	XCO uses K3s for management of services. These certificates are for secure communication of K3s with clients.
Default Validity Period	Expires in 10 years from the date of installation.
Impact on the system	
Renewal Procedure	K3s CA on page 57.
Alarm/Notification	Notification is sent to XCO subscribers: 1. After 30 days of expiry 2. Expired certs 3. Renewal certs

Intermediate CA Certificate of XCO

Location in TPVM deployment	/apps/efadata/certs/ca/extreme-ca-cert.pem
Location in server deployment	/opt/efadata/certs/ca/extreme-ca-cert.pem
Description	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
Default Validity Period	Expires in 10 years from the date of installation
Impact on the system	
Renewal Procedure	XCO Intermediate CA on page 53
Alarm/Notification	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

Location in TPVM deployment	/apps/efadata/certs/ca/extreme-ca-root.pem
Location in server deployment	/opt/efadata/certs/ca/extreme-ca-root.pem
Description	The certificate of Certificate Authority, which is the issuer of Intermediate CA certificate
Default Validity Period	Expires in 20 years from the date of installation
Impact on the system	
Renewal Procedure	XCO Root CA on page 52
Alarm/Notification	XCO Certificate Expiry Notice XCO Certificate Expired XCO Certificate Upload or Renewal

HTTPS Certificate of SLX

Location in TPVM deployment	/apps/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem
Location in server deployment	/opt/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem
Description	The certificate of SLX Web Server (Apache) for secure communication with the device from XCO
Default Validity Period	Expires in 2 years from installation
Impact on the system	System will not use encryption for HTTPS requests
Renewal Procedure	HTTPS Certificates on page 44
Alarm/Notification	Notification is sent to XCO subscribers from 30 days of expiry.

GRPC Certificate for OS ONE

Location in TPVM deployment	/apps/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem
Location in server deployment	/opt/efadata/certs/slx-<IP>.extremenetworks.com-cert.pem
Description	The certificate of OS ONE Web Server (Apache) for secure communication with the device from OS ONE
Default Validity Period	Expires in 2 years from installation
Impact on the system	The system will not use encryption for HTTPS requests
Renewal Procedure	HTTPS Certificates on page 44
Alarm/Notification	Notification is sent to XCO subscribers from 30 days before expiry.

K3s Certificate - XCO internal

Location in TPVM deployment	/apps/rancher/k3s/server/tls/
Location in server deployment	/var/lib/rancher/k3s/server/tls/
Description	XCO uses k3s for management of services. This certificate is for secure communication of k3s with clients
Default Validity Period	Expires in 1 year from installation. Reset after every upgrade of XCO
Impact on the system	
Renewal Procedure	K3s Server Certificate on page 58
Alarm/Notification	XCO Certificate Expiry Notice

JWT Signing or Verification - XCO internal

Location in TPVM deployment	/apps/efadata/certs/cert.crt.pem
Location in server deployment	/opt/efadata/certs/cert.crt.pem
Description	The RSA public key for JWT verification. This is also used to send user context from XCO to SLX. The same certificate is seen as Oauth certificate on SLX. It is applicable for both SLX and OS ONE.
Default Validity Period	Expires in 10 years from the date of installation
Impact on the system	Disables login to XCO
Renewal Procedure	JWT Certificate on page 57
Alarm/Notification	XCO Certificate Expiry Notice Managed Device Certificate Expiry Notice Managed Device Certificate Expired XCO Certificate Upload or Renewal Managed Device Certificate Upload or Renewal

Galera Certificate

Location in TPVM deployment	/apps/efadata/galera/galera.pem
Location in server deployment	/opt/efadata/galera/galera.pem
Description	The certificate enables SSL for the replication across the nodes. This is only applicable for multi-node deployment of XCO.
Default Validity Period	Expires in three years from the date of installation which is reset on every upgrade. There is no down time when the certificates are renewed.
Impact on the system	Replication of data between the nodes will fail.
Renewal Procedure	Galera Certificate on page 59
Alarm/Notification	NA

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 922.

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 922.

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 44 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 53 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

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
      --client-key string    Client Key File for the certificate
      --client-cert string   Client Certificate for EFA
```

Example:

```
efa certificate server --certificate=my_server.pem --key=my_server.key --cacert=ca-
chain.pem --client-cert=my_client.pem --client-key=my_client.key

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 922.

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 44.
- 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 922.

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 44 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 922.

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 922.

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.

Galera Certificate

XCO uses SSL encryption for the communication between the nodes in a multi-node deployment.

XCO uses the MariaDB and Galera services to implement an HA deployment.

- MariaDB and Galera – The registrations and configurations for devices, fabrics, and tenants in XCO are stored in a group of databases that are managed by MariaDB. Two nodes operate on a single MariaDB server and utilize Galera clustering technology to synchronize the business state in between the two nodes during standard operation.

When deploying the XCO cluster with multiple nodes, Galera components are automatically configured to communicate over SSL. This SSL configuration does not affect the communication between the cluster servers and their clients. During installation, the SSL configuration generates certificates for the Galera servers. These certificates are signed by the XCO Intermediate CA certificate and remain valid for three years from the date of installation. Upgrades reset the certificate validity period, and there's no downtime when renewing these certificates.

Location

- **TPVM:** `/apps/efadata/galera/galera.pem`
- **Server:** `/opt/efadata/galera/galera.pem`

Renewal

To renew Galera certificates, use the **`efa certificate server renew`** command.

For information about commands and supported parameters to renew Galera certificates, see .*ExtremeCloud Orchestrator Command Reference, 4.1.0*

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

Alarms for Auto Certificate Renewal Failure

XCO 3.8.0 introduces alerts and alarms for certificate renewal failures in addition to the existing certificate expiry notifications and certificate expired alerts and alarms. The new functionality doesn't impact existing certificate expiry alarms and alerts.

XCO currently has basic certificate management with manual enrollment and renewal, certificate validation, and expiration alarms. This feature expands on that by adding automated certificate renewal and integration with external Certificate Authorities (CAs).

Note
Renewal failures will trigger alerts and alarms, provided the system is functioning properly.

- Automatic Renewal Failure Alerts: New alarms when automatic renewal of internal certificates fails. This excludes automatic renewals during XCO upgrades, as the system won't be active.
- Manual Renewal of XCO CA Certificates: A new alert to capture renewal failures for specific certificates, including:
  - App Server Certificate
  - Default Intermediate CA
  - Default Root CA
  - WT Certificate
  - K3s Server Certificate

- K3s CA
- Galera Certificate

The following table describes the existing renewal mode for various certificate types:

Certificate Name	Validity	Existing Renewal Mode	CLI Command for Manual Renewal
App Server Certificate	3 yrs	Automatic Renewal during XCO upgrade and Manual	efa certificate server renew –cert-type server
Default Intermediate CA	10 yrs	Manual Renewal	efa certificate server renew –cert-type intermediate-ca
Default Root CA	20 yrs	Manual Renewal	efa certificate server renew –cert-type root-ca
JWT Certificate	10 yrs	Manual Renewal	efa certificate server renew –cert-type token
K3s Server Certificate	1 yrs	Automatic Renewal during XCO upgrade and manual	efa certificate server renew –cert-type k3s-server
K3s CA	10 yrs	Automatic Renewal during XCO upgrade and manual	efa certificate server renew –cert-type k3s-ca
Galera Certificate	3 yrs	Automatic Renewal during XCO upgrade and manual	efa certificate server renew –cert-type galera

Alert Details

```
{
    "Name": "CertificateRenewalFailureAlert",
    "AlertID": 31009,
    "Severity": "Warning",
    "Resource": "/App/System/Security/Certificate?type=*",
    "DisplayName": "Certificate renewal failure alert"
}
```

Alarm Details

```
{
        "Name": "CertificateRenewal",
        "AlarmID": 32002,
        "Type": "Security",
        "Severity": "Warning, Critical",
        "Resource": "/App/System/Security/Certificate?device_ip=*&type=*",
        "WillClear": true,
        "Description": "Notify when a certificate renewal is failed.",
        "RaiseCondition": [
            {
                "name": "CertificateRenewal_Raise",
                "desc": "Certificate Renewal has failed",
                "salience": 0,
                "when": "Alert.AlertID == 31009",
```

```
                "then": [
                    "Alert.Log(CertificateRenewal _Raise')",
                    "Alert.RaiseAlarm()",
                    "Retract('CertificateRenewal _Raise')"
                ]
            },
        "ClearCondition": [
            {
                "name": "CertificateRenewal_Clear",
                "desc": "Clear the CertificateRenewalFailure alarm when an application
certificate has been renewed.",
                "salience": 10,
                "when": "Alert.AlertID == 31004",
                "then": [
                    "Alert.Log('CertificateRenewal_Clear')",
                    "Alert.ClearAlarm()",
                    "Retract('CertificateRenewal_Clear')"
                ]
            }
        ]
    }
```

Certificate Management Using CMPv2

Learn about how to configure and manage certificates in XCO using CMPv2 (Certificate Management Protocol v2).

The CMPv2-based automated certificate lifecycle management feature introduced in XCO 4.1.0. The feature enables XCO to enroll and automatically renew northbound certificates through an external Certificate Authority (CA), reducing manual administrative overhead and the operational risk associated with certificate expiry.

XCO acts as a CMPv2 client, communicating with a customer-provided external CA (EJBCA) to request, install, and renew certificates for its northbound interfaces. Support for mutual TLS (mTLS) on RELP-based syslog connections is also included.

Use this feature to:

- Reduce manual certificate handling
- Improve operational reliability
- Maintain secure communication across services

Platforms Supported

- XCO 4.1.0
- Single-node and HA deployments

Deliverables

- CMPv2 certificate enrollment (Initial Request) and renewal (Key Update Request), supported on fresh deployments and on all supported upgrade paths from earlier XCO releases.
- CLI-based certificate lifecycle management, including apply, renew, status, and cleanup operations.

- RELP syslog integration with mutual TLS (mTLS) using CMPv2-provisioned client certificates.

Limitations

- Applying updated certificates may require a service restart or configuration reload to take effect.
- CMPv2 cannot be configured during installation. Enable it only after the system is fully operational.
- Only one external CA is supported per system.
- CMPv2 configuration is CLI-only in this release. UI workflows are not supported.
- CMPv2 and third-party certificate (TPC) management are mutually exclusive. Only one mechanism can be active at a time.
- CMPv2 certificates are not included in the XCO backup archives. A restore operation preserves the certificate management state active at the time of backup.
- CMPv2 certificates and related configuration are not included in XCO backups. When you restore an XCO backup, any existing CMPv2 configuration on the target system is retained. After the restore operation completes, you must reconfigure CMPv2. For best results, perform the restore on a newly deployed system.
- Before reverting to a different certificate management mechanism, run **`efa certificate cmpv2 clean`** to remove the active CMPv2 configuration.

Resilience and Redundancy

If certificate renewal fails, the existing certificate remains active and service continuity is maintained. An alert is generated on failure, and manual remediation is available through the CLI.

Backup and restore operations preserve the existing certificate and its associated configuration. If the system's IP address, VIP, or OVA IP changes, certificate regeneration may be required to reflect the updated identity.

Note
CMPv2 renewal must be initiated before certificate expiry. If a certificate has already expired, automatic renewal is not supported. To recover, reset the end-entity status to NEW on the external CA, then re-enroll using **`efa certificate cmpv2 apply`**.

To deregister CMPv2, run `efa certificate cmpv2 clean`. This removes all CMPv2 certificate artifacts and returns the system to a state where it can be reconfigured with XCO self-signed or third-party certificates.

Event Log Messages

CMPv2 enrollment, renewal, and failure events are captured in the following logs:

- Installer logs
- Monitor service logs

Include both sources when collecting diagnostic data for support cases.

Licensing

- EJBCA Community Edition is supported by default.
- EJBCA Enterprise Edition features, if required, must be licensed and provided by the customer.
- This feature introduces no additional XCO licensing requirements.

Performance and Scalability

- Renewal checks run once every 24 hours.
- Certificate validation is performed locally; no external validation calls are made during normal operation.
- CPU and memory impact is minimal.
- Expected service interruption during certificate enrollment or renewal: approximately 1.5–2.5 minutes.

Third-Party and Open-Source Components

This feature uses the following third-party and open-source components:

Component	Version	Purpose	License
OpenSSL	v3.5 or later	CMPv2 client implementation for certificate enrollment and renewal (RFC 4210)	Apache License v2.0
EJBCA	Community Edition	External Certificate Authority for CMPv2-based certificate issuance	LGPL v2.1 (Community Edition)

Licensing and distribution

- No third-party or open-source components are redistributed outside the XCO software package.
- This feature introduces no per-device royalties or usage-based licensing obligations.
- Customers are responsible for procuring and managing any licenses required for external CA products, including EJBCA Enterprise Edition.

EJBCA CMPv2 Pre-Configuration Requirements

Learn about the mandatory pre-configuration requirements for the external Certificate Authority (EJBCA) that must be completed before enabling CMPv2 in XCO.

Requirement	Description
EJBCA instance deployed	Deploy an external EJBCA instance and verify it is reachable from XCO over the required network path.
CMP enabled	Enable CMP protocol support in the EJBCA server configuration.

Requirement	Description
CMP alias configured	Create a CMP alias in EJBCA and configure it for client-mode operation.
End-entity and certificate profiles	Before initiating CMP enrollment, configure the required End Entity and Certificate Profiles on the EJBCA server. Enable the `Clear Text Password Storage` option in the End Entity Profile to allow HMAC-based authentication. This setting enables EJBCA to access the enrollment password and use it to calculate the cryptographic Message Authentication Code (MAC) for CMP requests.
HMAC credentials provisioned	Set up HMAC-based authentication credentials (username and password) on EJBCA for use with the `--username` and `--password` parameters.
CA trust chain available	Download the CA trust chain certificate from EJBCA and place it on the XCO host before running the apply command.
RELP mTLS: multiple certificates per entity	If RELP mTLS is required, configure the EJBCA end-entity profile to allow issuance of more than one certificate per entity. A profile restricted to a single certificate prevents the RELP client certificate from being provisioned.
RELP server trust	Install the CA chain (the same file passed to `--cacert` during apply) on the external RELP syslog server so it can validate XCO client certificates.

Configure CMPv2

You can configure CMPv2.

Before You Begin

Before you configure CMPv2, ensure the following:

- XCO 4.1.0 is installed and fully operational.
- An external EJBCA instance is deployed and reachable from XCO.
- CMP is enabled on the EJBCA server.
- A CMP alias is created and configured for client-mode operation.
- End-entity and certificate profiles for CMP enrollment are pre-configured on the EJBCA server.
- HMAC authentication credentials are provisioned on the external CA.
- The CA trust chain certificate is available locally on the XCO host.

About This Task

All CMPv2 certificate management operations are performed using the **`efa certificate cmpv2`** command set. Follow this procedure to configure CMPv2.

Procedure

1. Apply CMPv2 Certificate Management.

   The command registers an external CMPv2-compatible CA, enrolls certificates for XCO services, and optionally configures auto-renewal. This is the primary command for enabling CMPv2 certificate management.

   

   Important
   This is a disruptive operation. XCO services are unavailable for approximately 1.5–2.5 minutes while certificates are enrolled and applied. Do not run this command during a maintenance-sensitive period without prior planning.

```
efa certificate cmpv2 apply \
  --url <CMP-endpoint-URL> \
  --cacert <local-path-to-CA-chain> \
  [--cert <local-path>  --key <local-path>] \
  [--username <string>  --password <string>] \
  [--auto-renew enable|disable] \
  [--preferred-renewal-daytime <HH:MM:SS>]
```

   - Example 1 — Authentication with interactive password prompt (recommended):

```
$ efa certificate cmpv2 apply \
    --url https://10.37.12.53:8443/ejbca/publicweb/cmp/cmp-alias \
    --cacert MySubCA2-chain.pem \
    --username admin

Enter CMPv2 password:

CMPv2 certificate apply successful.
--- Time Elapsed: 2m25.041s ---
```

   - Example 2 — Authentication with auto-renewal enabled:

```
$ efa certificate cmpv2 apply \
    --cacert ManagementCA.crt \
    --username admin \
    --password <password> \
    --url https://10.37.11.38:8443/ejbca/publicweb/cmp/cmp-alias \
    --auto-renew enable

CMPv2 certificate apply successful.
--- Time Elapsed: 2m25.041s ---
```

   

   Caution
   Passing --password on the command line may expose credentials in shell history. Use the interactive prompt (omit --password) when possible.

2. View CMPv2 status.

   Displays the current CMPv2 configuration and the operational status of certificate enrollment and renewal.

```
$ efa certificate cmpv2 status
```

   Example Output — Certificate Installing/Installed/Renewing/Renewed

```
$ efa certificate cmpv2 status
+------+------------------+----------+----------+------------------------+------------|
| Name |       URL        | Username |  Status  |    Updated At          | Auto Renew
|
+------+------------------+----------+----------+------------------------+------------
+
```

```
| test | <CMP-Alias_URL>  | Admin    |Installing| 2026-04-10 09:27:09.83 |   true
|
+------+------------------+----------+----------+----------+-------------+-------------+
CMPv2 Certificate status detail
--- Time Elapsed: 40.6526ms ---

$ efa certificate cmpv2 status
+------+------------------+----------+----------+-------------------------+-------------|
| Name |       URL        | Username |  Status  |    Updated At           | Auto Renew
|
+------+------------------+----------+----------+-------------------------+-------------
+
| test | <CMP-Alias_URL>  | Admin    | Installed| 2026-04-10 09:27:09.83 |   true
|
+------+------------------+----------+----------+----------+-------------+-------------+
CMPv2 Certificate status detail
--- Time Elapsed: 40.6526ms ---

$ efa certificate cmpv2 status
+------+------------------+----------+----------+-------------------------+-------------|
| Name |       URL        | Username |  Status  |    Updated At           | Auto
Renew |
+------+------------------+----------+----------+-------------------------+-------------
+
| test | <CMP-Alias_URL>  | Admin    | Renewing | 2026-04-10 09:27:09.83 |   true
|
+------+------------------+----------+----------+----------+-------------+-------------+
CMPv2 Certificate status detail
--- Time Elapsed: 40.6526ms ---


$ efa certificate cmpv2 status
+------+------------------+----------+----------+-------------------------+-------------|
| Name |       URL        | Username |  Status  |    Updated At           | Auto
Renew |
+------+------------------+----------+----------+-------------------------+-------------
+
| test | <CMP-Alias_URL>  | Admin    | Renewed  | 2026-04-10 09:27:09.83 |   true
|
+------+------------------+----------+----------+----------+-------------+-------------+
CMPv2 Certificate status detail
--- Time Elapsed: 40.6526ms ---
```

3. Run manual certificate renewal.

   Initiates an immediate certificate renewal (KUR) request to the external CA. Use this command when auto-renewal is disabled or when renewal must be performed ahead of schedule.

   

   Caution
   Manual renewal causes a brief service disruption of approximately 1.5–2.5 minutes. A confirmation prompt is presented before the operation proceeds.

```
$ efa certificate cmpv2 renew
```

   Example Output:

```
$ efa certificate cmpv2 renew

WARNING: Executing above command may cause temporary downtime or service disruption.
Do you want to proceed? (yes/no): yes

Successfully renewed CMPv2 certificate.
```

4. Clean Up CMPv2 Configuration.

   Removes CMPv2 certificate artifacts and configuration from XCO, returning the system to its previous certificate management state. Run this command before switching to a different certificate management mechanism.

```
$ efa certificate cmpv2 clean
```

   Example Output:

```
$ efa certificate cmpv2 clean

CMPv2 certificate clean successful.
--- Time Elapsed: 2m43.084s ---
```

   Note
   After running efa certificate cmpv2 clean, also revoke or remove the corresponding end entity and certificate profile from the EJBCA server to complete cleanup on the CA side.

5. View Certificate Expiry and Managing Module.

   Displays the expiry date and managing module (CMPv2 or XCO Self Signed) for all certificates on the system. Use this command to verify which certificates are actively managed by CMPv2.

```
$ efa certificate expiry show
```

   Example Output:

```
$ efa certificate expiry show
+--------------------------------+--------+-----------------+----------------------+
| Name                           | Status | Managing Module | Expiry               |
+--------------------------------+--------+-----------------+----------------------+
| App Server Certificate         | Valid  | CMPv2           | 2026-05-05T09:37:45Z |
| Default Intermediate CA        | Valid  | XCO Self Signed | 2036-05-01T09:37:00Z |
| Default Root CA                | Valid  | XCO Self Signed | 2046-04-29T09:36:59Z |
| JWT Certificate                | Valid  | XCO Self Signed | 2036-05-01T09:40:53Z |
| K3s Server Certificate         | Valid  | XCO Self Signed | 2027-05-04T09:18:01Z |
| K3s CA                         | Valid  | XCO Self Signed | 2036-05-01T09:18:01Z |
| RELP Syslog client certificate | Valid  | CMPv2           | 2026-05-05T09:37:46Z |
| Galera Certificate             | Valid  | XCO Self Signed | 2036-04-07T07:09:55Z |
+--------------------------------+--------+-----------------+----------------------+
Certificate Expiry List detail show
--- Time Elapsed: 224.1ms ---
```

   Note
   Certificates showing 'CMPv2' in the Managing Module column are enrolled and renewed automatically. Certificates showing 'XCO Self Signed' continue to be managed by XCO internally and are not affected by CMPv2 operations.

Auto Renewal of Certificate

Enabled by default and triggered before expiration.

When auto-renewal is enabled, XCO evaluates certificate validity once every 24 hours. Renewal is triggered approximately seven days before the certificate expiry date. The renewal process

runs automatically in the background and does not require manual intervention under normal conditions.

Behaviour	Description
Renewal check frequency	Once every 24 hours.
Renewal window	Initiated approximately seven days before certificate expiry.
Failure handling	Existing certificates remain active if renewal fails. Alerts are generated and manual remediation is possible via CLI.
Expired certificate	Auto-renewal is not supported for already-expired certificates. See Section 5.1 for recovery steps.
IP or VIP address changes	Certificate regeneration may be required after IP address, VIP, or OVA IP changes to ensure the updated identity is reflected.

Switch Certificate Management Mechanism

You can switch from CMPv2 to TPC or XCO self-signed certificates.

About This Task

CMPv2 and third-party certificate (TPC) management are mutually exclusive.

Follow this procedure to switch from CMPv2 to TPC or XCO self-signed certificates. The transition requires explicit user confirmation and does not affect checkpointed system state.

Procedure

1. Remove CMPv2 configuration.

```
$ efa certificate cmpv2 clean
```

2. Apply the target certificate mechanism.
   a. To revert to XCO self-signed certificates, run the following command:

```
$ sudo mv $EFA_DATA_DIR/certs/cmpv2 /tmp/cmpv2_bkp
$ k3s kubectl delete secret efasecret-tls -n kube-system
$ efa certificate server renew --cert-type=server
```

   b. To apply third-party certificates, run the following command:

```
$ sudo mv $EFA_DATA_DIR/certs/cmpv2 /tmp/cmpv2_bkp
$ k3s kubectl delete secret efasecret-tls -n kube-system
$ efa certificate server --key=</path/to/key> \
    --certificate=</path/to/cert> \
    --cacert=</path/to/ca>
```

Note

If third-party certificates were previously installed, replace the first step with: `sudo mv $EFA_DATA_DIR/certs/thirdparty /tmp/thirdparty_bkp`.

Compatibility and Upgrade Considerations

Backward Compatibility

This feature does not introduce changes to existing checkpointed data structures and does not modify ISSU or rolling upgrade workflows. Backward compatibility is fully preserved between XCO 4.0.2 and XCO 4.1.0.

Certificate Mechanism Exclusivity

- CMPv2 and third-party certificate (TPC) management are mutually exclusive. Only one mechanism can be active at a time.
- CMPv2 takes precedence when configured.
- Transitioning between mechanisms requires running efa certificate cmpv2 clean before applying an alternative configuration.

Upgrade Path

CMPv2 certificate management is supported on fresh deployments and on supported upgrade paths from previous XCO releases. After upgrade, existing certificates remain active. CMPv2 must be explicitly enabled post-upgrade — it is not activated automatically.

Backup and Restore

CMPv2 certificates are not included in XCO backup archives. After a restore operation, the system preserves the certificate management state that was active prior to the restore. Re-enrollment through efa certificate cmpv2 apply may be required if the restored state does not include a valid CMPv2 configuration.

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.

OS ONE Certificate Management

The OS ONE device facilitates the importing of app certificates to establish a secure TLS connection. When a OS ONE device is either registered or discovered in the XCO, a unique certificate and key specific to the device will be generated. This certificate or key will be installed on the device using the CLI to enable communication between the device and XCO.

By default, the HTTPS port (443) is disabled, and there is no pre-configured gRPC server on the OS ONE. During the device registration process, the XCO will configure the gRPC server settings and activate the TCP port for the gRPC server to listen on (443). OS ONE provides certificate management capabilities through gNSI and the CLI. However, the XCO utilizes the CLI for certificate management.

The following certificates are installed when registering a OS ONE device:

- GRPC server certificate: A unique certificate for the device that facilitates secure communication with the gRPC server.
- Syslog CA certificate: Used for sending Syslog messages to the XCO over a secure port (6514)
- JWT token verifier certificate

Upon upgrading the OS ONE device, the existing certificates are kept. The default copy to startup command will also preserve the installed certificates.

Note
During the backup and restore process of the XCO, the new certificate must be installed on the device during its registration.

Configure gRPC Server Certificate

About This Task
Complete the following workflow to configure a gRPC server certificate:

Procedure

1. Generate Certificates on page 73.
2. Install gRPC Certificate on page 74.
3. Configure gRPC Server on page 76.
4. View Certificate Expiry on page 74.
5. Renew or Install gRPC Certificate on page 77.
6. Install Third-Party Certificate on page 77.
7. Configure App Certificate Alerts on page 78.
8. Remove or Uninstall App Certificate and gRPC Server Configuration on page 79.

Generate Certificates

The XCO generates and installs a device-specific certificate on the OS ONE during device registration, which is valid for three years. The pem certificate and key are created using the default CA certificate (extreme-ca-cert.pem) and key (extreme-ca-key.pem).

The gRPC certificate and key use the following naming convention: `oneos-<device IP>.extremenetworks.com-<cert|key>.pem`

The following is a sample certificate and key:
`oneos-10.64.204.43.extremenetworks.com-key.pem`

View Certificate Expiry

You can view the certificates expiry information.

About This Task
XCO fetches OS ONE certificates and compares them with its internal certificates. The response displays certificate type, expiration, and sync status, which can be "In Sync", "Not In Sync", or "Sync comparison error".

Follow this procedure to view the certificate expiry information.

Procedure

To view the device certificate expiry information, run the following XCO commands:

```
efa certificate device expiry show --ip <device IP address>
efa certificate device expiry show --fabric <fabric name>
efa certificate device expiry show -fabric-all
```

The following is a example output:

```
2028-04-23 04:15:30 +0000 GMT |
(efa:user)user@efa:~/xco/deploy/efa$ efa certificate device expiry show -- ip 10.64.188.63
+--------------+-----------+-------------------------------+------------------------
+----------+
|      IP      |   Type    |             Expiry            | EFA/Device
Sync Status | Status  |
+--------------+-----------+-------------------------------+------------------------
+----------+
| 10.64.188.63 |   GRPC    | 2028-04-23 04:15:30 +0000 GMT |
In Sync        | Success |
+--------------+-----------+-------------------------------+------------------------
+----------+
| 10.64.188.63 | SyslogCA  | 2045-05-18 04:02:40 +0000 UTC |
In Sync        | Success |
+--------------+-----------+-------------------------------+------------------------
+----------+
```

Install gRPC Certificate

You can install or import a gRPC certificate.

About This Task
The XCO runs a CLI command on the OS ONE device to import the gRPC certificate and corresponding key from a remote host. Once imported, the certificate will be associated with an SSL profile ID, which is then used in the gRPC server configuration.

Note
XCO uses a fixed SSL profile ID: `xco_grpc_ssl_profile_id`. If this profile ID already exists on the device, it will be overwritten during the import process.

Follow this procedure to install or import the gRPC certificate.

Procedure

1. To import the certificate, run the following command:

   Note
   XCO currently supports only the `scp` protocol for importing certificates.

   For information on command syntax and parameter description, see *ExtremeCloud Orchestrator Command Reference Guide*.

```
certificate-manager import ssl-profile-id certificate-id app-cert protocol {scp |
sftp} host host-address certificate-pem-file key key-pem-file user username password
password
```

   The following is an example command output:

```
certificate-manager import ssl-profile-id xco_grpc_ssl_profile_id
app-cert protocol scp host 10.32.85.231 certificate /opt/
efadata/certs/tos-10.64.188.63.extremenetworks.com-cert.pem key /opt/efadata/certs/
tos-10.64.188.63.extremenetworks.com-key.pem user efainternal password ************

Output: certificate-manager import response: Imported xco_grpc_ssl_profile_id
certificate successfully
```

2. View imported certificates on the device.
   a. View Application Certificates.
      i. To list all application certificates installed on the device, run the following command:

```
show certificate-manager app-cert all
```

      ii. To view a specific application certificate by SSL profile ID, run the following command:

```
show certificate-manager app-cert <ssl_profile_id>
```

      The

      The following is an example command output:

```
vtos63# show certificate-manager app-cert all
App level certificates:
certificate-id:xco_grpc_ssl_profile_id
Endpoints using this certifcate-id: [type: EP_DAEMON endpoint: "grpc-server
xco_grpc_server_443"]
sha256 Fingerprint-BC: 4A:D5:C9:2E:0A:
EA:24:CD:20:3D:F1:99:32:13:5E:4D:72:5C:6D:33:87: AC:EE:DB:E9:66:80:2E:05:04:BE
X509v3 Subject Alternative Name:
DNS: tos-10.64.188.63.extremenetworks.com, IP Address:10.64.188.63
subject=CN=tos-10.64.188.63.extremenetworks.com
issuer=C=US, ST=CA, O=Extreme Networks, OU=Extreme Fabric Automation Intermediate,
CN=EFA Intermediate CA, emailAddress=support@extremenetworks.com notBefore=Apr 19
10:35:15 2025 GMT
notAfter-Apr 19 10:35:15 2028 GMT
```

   b. View CA Certificates.
      i. To list all imported CA certificates, run the following command:

```
show certificate-manager ca-certificates all
```

      ii. To filter and view a specific CA certificate by SSL profile ID, run the following command:

```
show certificate-manager ca-certificates <ssl_profile_id>
```

      The following is an example command output:

```
vtos63# show certificate-manager ca-certificates all
CA certificates:
certificate-id: xco_ca_ssl_profile_id
Endpoints using this certifcate-id: [type: EP_DAEMON endpoint: "logging remote-
```

```
server 10.32.85.231"]
sha256 FingerPrint-DB:3F:95:8E:68:76: C2: BF:43:DF:BA:48:4F:29:68: E1:77: A1:5A:
3D:7F:11:B3:13:E5:33:62:81: FC: BA:ED:56
56e6574776f726b732e636f6d
subject=CN=EFA Intermediate CA, OU=Extreme Fabric
Automation Intermediate, 0=Extreme Networks,
ST=CA, C=US, 1.2.840.113549.1.9.1=#0c1b737570706f72744065787472656d6
issuer=CN=efa.extremenetworks.com,OU=Extreme Fabric Automation, 0=Extreme Networks,
L=SJ, ST=CA, C=US, 1.2.840.113549.1.9.1=#0c1b737570706f72744065787472656d656e65
74776f726b732e636f6d
notBefore=May 17 07:38:21 2025 UTC
notAfter May 15 07:38:21 2035 UTC
sha256 FingerPrint=F3:26:BA:C6:4A:25:A3:B4:DC: 83:5A: FC:93:E4:50:94:B6:C6:7E:
AC:8C:F3:07:B0:01:80:33:33:3D:B2:DF:F0
574776f726b732e636f6d
subject=CN=efa.extremenetworks.com, OU=Extreme Fabric Automation,
0=Extreme Networks, L=SJ, ST=CA, C=US, 1.2.840.113549.1.9.1-
#0c1b737570706f72744065787472656d656e6 issuer=CN=efa.extremenetworks.com,OU=Extreme
Fabric Automation, 0=Extreme Networks, L=SJ, ST=CA, C=US,
1.2.840.113549.1.9.1=#0c1b737570706f72744065787472656d656e65 notBefore=May 17
07:38:17 2025 UTC
74776f726b732e636f6d
notAfter=May 12 07:38:17 2045 UTC
```

Configure gRPC Server

You can configure a gRPC server.

About This Task

OS ONE does not create a default gRPC server. Instead, XCO configures the gRPC server, which includes setting the name, port, and certificate ID (SSL profile ID used for the gRPC certificate). The gRPC server name used by XCO is "xco_grpc_server_443".

Follow this procedure to configure a gRPC server.

Procedure

1. To configure a gRPC server, run the following command:

```
grpc-server <Name of the gRPC instance>
    certificate-id <Name of the certificate that is associated with the gRPC service>
    enable <Enable northbound access>
    port [1-65535] <TCP port on which the gRPC server should listen>
```

If a gRPC server with the specified name already exists during device registration, it will be removed before configuring the new gRPC server. The certificate ID must match the SSL profile ID used when importing the gRPC certificate to the device.

XCO returns an error if the specified ports and VRF are already configured with another server configuration. XCO only manages the gRPC server configuration with the name "xco_grpc_server_443".

The following is an example of the gRPC server configuration:

```
grpc-server xco_grpc_server_443
    certificate-id xco_grpc_certificate_id
    port 443
    enable
```

2. Verify the gRPC server configuration.

```
show running-config system grpc-server
```

The following is an example output:

```
vtos63# show running-config system grpc-server system
  grpc-server xco_grpc_server_443
    certificate-id xco_grpc_ssl_profile_id
    port 443
    enable
```

Renew or Install gRPC Certificate

You can renew or install a gRPC certificate.

About This Task

You can manually regenerate and reinstall a certificate on a device using the EFA command below. This command will remove the existing certificate and generate and install a new one on the specified device(s). If a valid certificate is already present, the command will be skipped by default—unless the `--force` option is used.

Follow this procedure to renew or install a gRPC certificate

.

Procedure

To renew or install a gRPC certificate, run the following command:

```
efa certificate device install [ [ --ip ip-addr | --fabric-all | --fabric fabric-name]
--certType { https | token | grpc } | -https-certificate --https-key | -grpc-certificate
--grpc-key  --force]
```

The following is an example CLI output:

```
efa certificate device install --ip 10.x.x.x  --certType=grpc --force
efa certificate device install --fabric fb
efa certificate device install --fabric-all
efa certificate device install --ip
x.x.x.x --grpc-certificate=tos-10.64.216.44.extremenetworks.com-cert.pem --grpc-
key=tos-10.64.216.44.extremenetworks.com-key.pem
```

When you run this command, XCO will trigger the **`certificate-manager import`** command to rotate the certificate on the OS ONE device. After renewal, the certificate ID is re-associated with the gRPC server configuration. This process involves disabling and re-enabling the gRPC server to apply the new certificate.

If the base CA files (for example, `root-ca`, `intermediate-ca`) are renewed, run this command to generate and install the updated gRPC certificate on the devices.

Install Third-Party Certificate

You can install a third-party certificate.

About This Task

To import a third-party gRPC certificate, use the 'grpc-certificate' and 'grpc-key' parameters with the PEM file and key. To revert to XCO certificates, use the same command and select the XCO options. Third-party certificates are stored in the /opt/inventory/slxcerts/thirdparty (host path /opt/efadata/certs/thirdparty) directory.

Follow this procedure to install a third-party certificate.

Procedure

To install a third-party certificate, run the following command:

```
efa certificate device install --ip 10.x.x.x --cert-type grpc --grpc-certificate
device190.pem --grpc-key device190.key
```

Results

Configure App Certificate Alerts

You can configure an app certificate alerts.

About This Task

The gRPC certificate generated by XCO is valid for three years. OS ONE devices support certificate expiry alerts through RASlog messages, providing advanced notifications based on severity levels and the number of days remaining until certificate expiration.

The certificate expiry alert feature in OS ONE supports multiple severity levels: **critical**, **major**, **minor**, and **info**. Each severity level can be configured with a custom alert threshold (in days) before the certificate expires. The configurable range is from 1 to 90 days, and setting the value to 0 disables the alert for that level.

Follow this procedure to configure app certificate alerts.

Procedure

1. To configure certificate expiry alerts using the EFA CLI, run the following commands:

   ```
   efa inventory device setting update --crypto-cert-expiry-info <time period>
   efa inventory device setting update --crypto-cert-expiry-minor <time period>
   efa inventory device setting update --crypto-cert-expiry-major <time period>
   efa inventory device setting update --crypto-cert-expiry-critical<time period>
   ```

   Replace the `<time period>` with the desired number of days for each severity level.
2. Run the following device CLI on OS ONE:

   ```
   Device43(config-system-cert-mgr-exp)# info 90
   Device43(config-system-cert-mgr-exp)# minor 60
   Device43(config-system-cert-mgr-exp)# major 30
   Device43(config-system-cert-mgr-exp)# critical 10
   ```

Example

The following is an example CLI output that configures the certificate expiry based on the severity and number of days.

```
efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiryinfo
50

efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-
expirycritical 10
```

Remove or Uninstall App Certificate and gRPC Server Configuration

You can remove or uninstall an app certificate and a gRPC server configuration.

About This Task

When the OS ONE device is unregistered from XCO, its gRPC server configuration and certificate are retained because XCO cannot confirm configuration cleanup across services. If the device is re-registered, the existing configuration and certificate will be cleared before applying the new settings.

Follow this procedure to remove or uninstall an app certificate and gRPC server configuration.

Procedure

1. To remove the certificates from the device, run the following command:

```
certificate-manager delete ssl-profile-id all
```

2. To remove a specific certificate from the device, provide the corresponding SSL profile ID.

```
certificate-manager delete ssl-profile-id <profile-id>
```

Configure Syslog CA Certificate

About This Task
Complete the following workflow to configure a Syslog CA certificate:

Procedure

1. Install Syslog CA Certificate on page 79
2. Configure Syslog Server on page 79
3. Syslog CA Alerts on page 80
4. Syslog CA Certificate Uninstallation

Install Syslog CA Certificate

Syslog will be configured automatically during device registration.

About This Task

During device registration, the inventory service runs the **`import ca-cert`** command to install the default intermediate CA certificate on the device. Upon successful installation of the certificate, the secured Syslog server (XCO) is configured on the device.

Configure Syslog Server

You can configure a syslog server.

About This Task
XCO configures a secured Syslog server for the OS ONE device after installing the Intermediate-CA certificate. XCO will periodically verify the status of the Syslog server.

Follow this procedure to configure a syslog server.

Procedure

To configure a syslog server, run the following command:

```
kvm-x86-64(config-system-logging-remote-server-h1)# source-address 1.1.1.1
Warning: Existing Host configuration changed
kvm-x86-64(config-system-logging-remote-server-h1)# remote-port
514-530)  Remote server port number <514-530> [default:514]
kvm-x86-64(config-system-logging-remote-server-h1)# remote-port 525
Warning: Existing Host configuration changed
kvm-x86-64(config-system-logging-remote-server-h1)# mode-transport       tcp  TCP
transport mode
kvm-x86-64(config-system-logging-remote-server-h1)# vrf                  NAME  <String:
1-64 character> vrf name
kvm-x86-64(config-system-logging-remote-server-h1)# tls-profile-id xco_ssl_profile_id
```

Example

The following is an example output of a syslog server configuration:

```
kvm-x86-64(config-system-logging-remote-server-h1)# tls  TLS Encryption (Use crypto
import command, to import CA certificate))
kvm-x86-64(config-system-logging-remote-server-h1)# secure-forwarding tls
kvm-x86-64(config-system-logging-remote-server-h1)# remote-port 6514
Warning: Existing Host configuration changed
kvm-x86-64(config-system-logging-remote-server-h1)# tls-profile-id
xco_ssl_profile_id                                                  kvm-
x86-64(config-system-logging-remote-server-h1)# mode-transport tcp
kvm-x86-64(config-system-logging-remote-server-h1)# vrf mgmt-vrf
```

Syslog CA Alerts

You can configure syslog CA alerts.

XCO regularly checks and validates the CA certificate on the device. It retrieves and validates trust store certificates and raises an alert if a certificate is missing or expired.

Syslog CA Certificate Uninstallation

Upon device deletion, XCO automatically purges both the Syslog configuration and the installed CA certificate from the device's runtime state.

Audit Logging and JWT Token Integration

Audit logging on the device enables tracking of users who log in and perform configuration changes by analyzing device logs. The username of the individual logged into XCO is transmitted to the device using the XCO access token.

The SSL profile feature is supported starting from version ExtremeOneSR-22.2.0.0-037 and later. During device registration, XCO sends a CLI command to import the `JWT verifier certificate (cert.crt.pem)` into the OS ONE device. An access token is generated for the logged-in user and included in the **gRPC client connection request**. If no access token is available, the client will fall back to using the **username and password** for authentication. The access token is valid for 24 hours.

When a device is deleted, XCO retains the J**WT SSL profile**. Upon re-registration of the same device, XCO first cleans up the existing SSL profile before importing the certificate again.

Configure Audit Logging

You can configure audit logging.

About This Task

Follow this procedure to configure audit logging.

Procedure

1. To import certificates, run the following device CLI commands:

```
certificate-manager import ssl-profile-id systemx-jwt-ssl-profile-id app-cert protocol
scp host 10.32.85.231 certificate /opt/efadata/certs/cert.crt.pem user efainternal
password password vrf mgmt-vrf
```

2. To associate the SSL profile with JWT Validator, run the following commands:

```
configure terminal
    system
        aaa
             token-validator systemx-jwt-token-validator
             ssl-profile-id systemx-jwt-ssl-profile-id
```

   The access token's CustomClaim now includes two custom properties:

   - **Role**: set to "admin"
   - **Requester**: set to "xco"

   A predefined SSL profile (`systemx-jwt-ssl-profile-id`) is linked to the imported certificate and tied to the JWT token validator.

Example

- Sample Output for Configuration Audit Logs

```
vtos63# show logging audit config | inc 0/5
2025-04-11 10:23:40.836 GMT +0000 LogID:6001 info Msg: user/10.32.85.231/http/grpc
Method:/gnmi.gNMI/Set Status:OK Request:prefix:{elem:{name:"interfaces"}} update:{path:
{elem:{name:"interface" key:{key:"name" value:"ethernet 0/5"}} elem:{name:"config"}
elem:{name:"description"}} val:{string_val:"testIntf"}}
```

- Sample Output for Security Audit Logs

```
vtos63# show logging audit security | inc 2025-04-11 10:23
2025-04-11 10:23:37.439 GMT +0000 LogID:7004 info Msg: Access token is generated for
the user: admin
2025-04-11 10:23:37.455 GMT +0000 LogID:7002 info Msg: User authentication is
successful for user: admin from 10.32.85.231
2025-04-11 10:23:40.552 GMT +0000 LogID:7002 info Msg: User authentication is
successful for user: user
2025-04-11 10:23:40.816 GMT +0000 LogID:7002 info Msg: User authentication is
successful for user: user
2025-04-11 10:23:41.103 GMT +0000 LogID:7002 info Msg: User authentication is
successful for user: user
```

Configure Device Security Settings

About This Task

Complete the following workflow to configure device security settings:

Procedure

1. View Security Settings on page 84
2. Minimum TLS Version on page 85
3. Copy Configuration on page 93
4. Factory Reset or Write Erase on page 93

Configure and View Security Settings

You can view the security setting and its status.

About This Task

Follow the procedure to configure and view the security settings and their status on a device.

Note

- Backup and Restore: Security settings are saved in the database and will be restored during backup and restore operations.
- Supportsave: All operation logs are included in the inventory support save and are preserved for troubleshooting and review.
- TPVM upgrade and HA multi-node: These feature will have no impact.

Procedure

1. To enable or disable global security settings on a device, run the following XCO command:

   The enable and disable commands determine whether global secure settings are applied to a device during registration. When enabled, these settings are pushed to the device; when disabled, they are not.

```
efa inventory device secure settings enable
secure settings enable[Success]
--- Time Elapsed: 43.222199ms ---

efa inventory device secure settings disable
secure settings disable[Success]
--- Time Elapsed: 53.016327ms --
```

2. After enabling secure settings, run the following XCO command to view the security setting status on a device:

   Note

   - OS ONE: `min-tls-version` defaults to 1.2, can be set to 1.3 (ignores 1.1).
   - SLX: Supports 1.1, 1.2, and 1.3.

```
efa inventory device secure settings show
+------------------------------+--------------------------------------+
|             Name             |                Value                 |
+------------------------------+--------------------------------------+
| Min-tls-version              | 1.3                                  |
+------------------------------+--------------------------------------+
| SSH Mac-algorithm            | hmac-sha2-512-etm@openssh.com        |
|                              | hmac-sha2-256-etm@openssh.com        |
|                              | umac-128-etm@openssh.com             |
|                              | hmac-sha2-256 hmac-sha2-512          |
+------------------------------+--------------------------------------+
| SSH Key-exchange-algorithm   | curve25519-sha256                    |
```

```
|                               | curve25519-sha256@libssh.org         |
|                               | diffie-hellman-group-exchange-sha256 |
|                               | diffie-hellman-group16-sha512        |
|                               | diffie-hellman-group18-sha512        |
|                               | diffie-hellman-group14-sha256        |
+-------------------------------+--------------------------------------+
| SSH Cipher                    | non-cbc                              |
+-------------------------------+--------------------------------------+
| Telnet                        | Disabled                             |
+-------------------------------+--------------------------------------+
| Max-password-age              | 90                                   |
+-------------------------------+--------------------------------------+
| Force-default-password-change | Disabled                             |
+-------------------------------+--------------------------------------+
| Secure-Settings               | Enabled                              |
+-------------------------------+--------------------------------------+
Common security settings
--- Time Elapsed: 174.180014ms ---
```

3. After disabling secure settings, run the following XCO command to view the security setting status on a device:

```
efa inventory device secure settings show
+-------------------------------+--------------------------------------+
|             Name              |                Value                 |
+-------------------------------+--------------------------------------+
| Min-tls-version               | 1.3                                  |
+-------------------------------+--------------------------------------+
| SSH Mac-algorithm             | hmac-sha2-512-etm@openssh.com        |
|                               | hmac-sha2-256-etm@openssh.com        |
|                               | umac-128-etm@openssh.com             |
|                               | hmac-sha2-256 hmac-sha2-512          |
+-------------------------------+--------------------------------------+
| SSH Key-exchange-algorithm    | curve25519-sha256                    |
|                               | curve25519-sha256@libssh.org         |
|                               | diffie-hellman-group-exchange-sha256 |
|                               | diffie-hellman-group16-sha512        |
|                               | diffie-hellman-group18-sha512        |
|                               | diffie-hellman-group14-sha256        |
+-------------------------------+--------------------------------------+
| SSH Cipher                    | non-cbc                              |
+-------------------------------+--------------------------------------+
| Telnet                        | Disabled                             |
+-------------------------------+--------------------------------------+
| Max-password-age              | 90                                   |
+-------------------------------+--------------------------------------+
| Force-default-password-change | Disabled                             |
+-------------------------------+--------------------------------------+
| Secure-Settings               | Disabled                             |
+-------------------------------+--------------------------------------+
Common security settings
--- Time Elapsed: 103.057551ms ---
```

4. To restore a specific flag under Secure Settings to its default value, run the following reset command:

```
efa inventory device secure settings reset --telnet
secure settings reset[Success]
--- Time Elapsed: 26.826033ms ---
```

5. To modify the values of the different flags available for the Secure Settings, run the following update command:

```
efa inventory device secure settings update --key-exchange-algorithm
curve25519sha256,curve25519-sha256@libssh.org
secure settings update[Success]--- Time Elapsed: 42.91008ms --
```

6. To push Secure Settings configurations to the device, run the following apply command. Until the apply command is run, the configurations won't take effect. You can apply settings to a single device using the `--ip` flag and to an entire fabric using the `--fabric flag`.

```
efa inventory device secure settings apply --ip 10.64.xxx.xx
+--------------+---------+---------+
|  IP Address  | Status  | Warning |
+--------------+---------+---------+
| 10.64.xxx.xx | Success |         |
+--------------+---------+---------+
Secure settings apply
--- Time Elapsed: 8.03212659s --
```

View Security Settings

You can view the security settings.

About This Task

Follow the procedure to view the security settings.

Procedure

To view the security settings on a device, run the following XCO command:

The 'show' command displays the global secure setting values stored in the DB. Use an `-ip` flag to check the device-level secure settings.

```
efa inventory device secure settings show -ip 10.x.x.x
+--------------------------------+----------------------------------------+
|              Name              |                 Value                  |
+--------------------------------+----------------------------------------+
| Min-tls-version                | 1.2                                    |
+--------------------------------+----------------------------------------+
| SSH Mac-algorithm              | hmac-sha2-512-etm@openssh.com          |
|                                | hmac-sha2-256-etm@openssh.com          |
|                                | umac-128-etm@openssh.com               |
|                                | hmac-sha2-256 hmac-sha2-512            |
+--------------------------------+----------------------------------------+
| SSH Key-exchange-algorithm     | curve25519-sha256                      |
|                                | curve25519-sha256@libssh.org           |
|                                | diffie-hellman-group-exchange-sha256 |
|                                | diffie-hellman-group16-sha512          |
|                                | diffie-hellman-group18-sha512          |
|                                | diffie-hellman-group14-sha256          |
+--------------------------------+----------------------------------------+
| SSH Cipher                     | non-cbc                                |
+--------------------------------+----------------------------------------+
| Telnet                         | Disabled                               |
+--------------------------------+----------------------------------------+
| Max-password-age               | 90                                     |
+--------------------------------+----------------------------------------+
| Force-default-password-change | Disabled                               |
+--------------------------------+----------------------------------------+
```

The enable and disable commands determine whether global secure settings are applied to a device during registration. When enabled, these settings are pushed to the device; when disabled, they are not.

Use the reset command to restore a specific flag under Secure Settings to its default value.

Use the update command to modify the values of the different flags available for the Secure Settings.

The apply command pushes Secure Settings configurations to the device. Until the apply command is run, the configurations won't take effect. You can apply settings to a single device using the `--ip` flag and to an entire fabric using the `--fabric` flag.

.

Minimum TLS Version

Learn about the supported platforms, TLS versions, services for override (OS ONE), configuration behaviour and validation.

The XCO mandates a minimum TLS version of 1.2. The Extreme OS ONE supports TLS versions 1.2 and higher, with 1.2 set as the default.

The TLS configuration feature enables you to:

- Enforce secure communication policies
- Centrally manage TLS versions across devices
- Customize TLS behavior for specific services (Extreme OS ONE)
- Configure a global minimum TLS version that applies to all TLS-based services on managed switches.
- Override the global setting per service on Extreme OS ONE, enabling individual services to enforce a different minimum TLS version.
- Manage all TLS settings through the CLI or REST API.
- Global settings apply automatically at device registration and persist across backup and restore operations — no reconfiguration required.

Use global settings for consistency and service-level overrides for flexibility where needed.

For a detailed procedure on configuring the minimum TLS version on devices, see Configure Minimum TLS Version on Device on page 90.

Supported Platforms

Platform	Global TLS Configuration	Service-Level Overrides	Automatic fallback if version unsupported	Reset to default
SLX-OS	Supported	Not supported	Supported	Supported
Extreme OS ONE	Supported	Supported	Supported	Supported

Global configuration applies across all platforms, while service-level overrides are available only on Extreme OS ONE.

Note

- Security settings are stored in the database and restored during backup/restore operations
- Logs are available through inventory service logs for troubleshooting
- No impact on HA, upgrades, or system performance
- Per-service TLS overrides are available on Extreme OS ONE only. Specifying `--min-service-tls-version` on an SLX device returns an error.

Supported TLS Versions

Version	Status	Notes
TLS 1.2	Supported — default	Applied when no global or service-level override is configured.
TLS 1.3	Supported	Can be set globally or per service. If a device does not support TLS 1.3, TLS 1.2 is applied automatically.

Supported Services for Override (OS ONE Only)

Service	CLI Identifier	TLS Use Case
gRPC Server	GRPC_SERVER	Secures gRPC-based internal API communication over TLS.
LDAP Client	LDAP_CLIENT	Secures LDAPS (LDAP over TLS) connections for user authentication.
RADIUS Client	RADIUS_CLIENT	Secures RADIUS over TLS (RadSec) for user authentication.
HTTPS Client	HTTPS_CLIENT	Secures outbound HTTPS traffic for firmware downloads and API calls.
Syslog Client	SYSLOG_CLIENT	Secures log forwarding to remote servers when TLS transport is enabled.

TLS Configuration Behavior Matrix — Extreme OS ONE

The following table shows the effective minimum TLS version applied to each service based on the combination of global and service-level settings. The same logic applies to all five supported services.

Service	Global TLS	Service TLS Set	Min-TLS Applied	Behavior
LDAP Client	1.2	Not set	1.2	Inherits global setting.
	1.2	1.3	1.3	Service override applied.
	1.3	Not set	1.3	Inherits global setting.
	1.3	1.2	1.2	Service override applied.
gRPC Server	1.2	Not set	1.2	Inherits global setting.
	1.2	1.3	1.3	Service override applied.
	1.3	Not set	1.3	Inherits global setting.
	1.3	1.2	1.2	Service override applied.
HTTPS Client	1.2	Not set	1.2	Inherits global setting.
	1.2	1.3	1.3	Service override applied.
	1.3	Not set	1.3	Inherits global setting.
	1.3	1.2	1.2	Service override applied.
RADIUS Client	1.2	Not set	1.2	Inherits global setting.
	1.2	1.3	1.3	Service override applied.
	1.3	Not set	1.3	Inherits global setting.
	1.3	1.2	1.2	Service override applied.
Syslog Client	1.2	Not set	1.2	Inherits global setting.
	1.2	1.3	1.3	Service override applied.

Service	Global TLS	Service TLS Set	Min-TLS Applied	Behavior
	1.3	Not set	1.3	Inherits global setting.
	1.3	1.2	1.2	Service override applied.

- Service-level configuration **overrides** global TLS settings
- If no service override is defined, the global setting is used
- Duplicate service entries are automatically consolidated
- If multiple versions are specified for a service, the **last value takes precedence**

Validation and Error Handling

The following table describes how the CLI handles edge cases, redundant entries, and unsupported configurations in the `--min-service-tls-version` parameter.

Scenario	Allowed	Action Taken	CLI / Log Behavior
Per-service TLS version equals global TLS version	**Yes**	No-op — configuration accepted but no change applied.	Success — no warning generated.
Same service listed multiple times with the same TLS version	**Yes**	Duplicate entries de-duplicated; single entry applied.	Success — no warning generated.
Same service listed with two different TLS versions	**Yes**	Last occurrence in the list takes precedence; earlier entries are overridden.	Success — final value applied.
Multiple valid unique service overrides (for example, LDAPv1.3, HTTPSv1.2)	**Yes**	All overrides applied as specified.	Success.
Service override specified for SLX (not supported)	**No**	Override rejected or ignored depending on context.	Error: service-level TLS not supported on SLX.
Invalid TLS version value (for example, LDAPv1.4)	**No**	Request rejected immediately.	Error: unsupported TLS version specified.

Note
When the same service appears multiple times with different versions in a single --min-service-tls-version value, the last occurrence wins. For example, LDAP_CLIENTv1.3,HTTPS_CLIENTv1.3,LDAP_CLIENTv1.2 results in LDAP Client using TLS 1.2 (the last value for that service).

Best Practices

- Use **TLS 1.3** wherever supported for enhanced security
- Apply service-level overrides only when required
- Validate configurations after applying changes
- Maintain consistent security policies across environments

Operational Considerations

- Backup and Restore
  - Security settings (TLS version, service overrides) are backed up and restored automatically with XCO database.
  - No manual reconfiguration needed after restore.
- Support Save
  - TLS config change logs are included in support save output.
  - Provide support save output for TLS-related support cases.
- TPVM Upgrade
  - TLS config is preserved during TPVM upgrades.
- HA Multi-Node
  - TLS config is synced across HA nodes automatically.
  - No extra steps needed for secondary nodes.
- Logging and Events
  - TLS config events are logged for auditing and troubleshooting.
  - Check inventory service logs for config changes and errors.

How TLS Version Configuration Works

XCO manages TLS version configuration at two levels: a global setting that applies to all TLS-based services on a device, and optional per-service overrides that allow individual services to enforce a different minimum version.

Global Setting

The global minimum TLS version sets the floor for all TLS negotiations on the device. XCO applies the current global setting automatically when a device is registered. Any subsequent change to the global setting is pushed to all registered devices immediately.

If a device does not support the configured version — for example, TLS 1.3 — XCO automatically applies the next supported version (TLS 1.2) for that device.

Per-Service Override (OS ONE Only)

On Extreme OS ONE platforms, you can override the global TLS setting for individual services. An override applies only to the specified service — all other services continue to use the global value. You can set the override to a value higher or lower than the global minimum.

Note
Per-service overrides are independent of each other. Changing the global setting does not remove or reset existing service-level overrides.

TLS Handshake Flow

XCO manages TLS configuration on the device. The following table describes how a client establishes a TLS session with an Extreme OS ONE device.

Caution
If a client proposes only TLS versions below the configured minimum, the handshake fails and the connection is rejected. Verify that all clients accessing managed services support the configured minimum TLS version before applying changes.

Step	Actor	Action
1	TLS Client	Sends Client Hello to the OS ONE device, proposing supported TLS versions and cipher suites.
2	OS ONE Device	Responds with Server Hello, selecting the highest mutually supported TLS version at or above the configured minimum.
3	Client + Device	Exchange cryptographic keys to establish a shared session secret
4	Client + Device	Begin encrypted data exchange over the established TLS session.5
5	XCO	Applies any pending configuration changes to the device as required.

Configure Minimum TLS Version on Device

You can configure the minimum TLS version on SLX-OS and OS ONE devices.

Before You Begin

Ensure the following:

- Devices are discovered and managed in XCO
- You have CLI access to the XCO environment
- Required privileges are available to update security settings

About This Task

The EFA CLI enables you to configure the minimum Transport Layer Security (TLS) version for managed devices. This ensures that all TLS-based communications meet defined security standards and prevents the use of outdated protocols.

You can configure:

- A **global minimum TLS version** for all services
- **Service-specific TLS overrides**

For information on commands and syntax for configuring the minimum supported TLS version in XCO and OS ONE, see *ExtremeCloud Orchestrator Command Reference Guide*.

Follow this procedure to configure the minimum TLS version on devices.

Procedure

1. Configure Global Minimum TLS Version.

   Set a minimum TLS version that applies to all supported services on the device.

   ```
   efa inventory device secure settings update --min-tls-version 1.3
   ```

   

   Note

   - If a device does not support the specified version, a lower supported version is applied automatically.
   - The configuration is applied to all applicable services by default.

2. Configure Service-Level TLS Overrides (OS ONE Only)

   Set the global minimum to TLS 1.2 and override LDAP Client and HTTPS Client to use TLS 1.3.

   

   Caution

   The `--min-service-tls-version` parameter is not supported on SLX. Using this flag against an SLX device returns an error.

   ```
   efa inventory device secure settings update
   --min-tls-version 1.2
   --min-service-tls-version LDAP_CLIENTv1.3,HTTPS_CLIENTv1.3
   ```

   

   Note

   - Override the global TLS setting for specific services.
   - Global TLS is set to **1.2**
   - Specified services use **TLS 1.3**
   - Service-level configuration overrides the global setting

3. View Security Settings.

   a. View Global security settings.

   ```
   efa inventory device secure settings show
   ```

   Example:

   ```
   $ efa inventory device secure settings show

   +-------------------------------+---------------------------------------+
   |             Name              |                Value                  |
   +-------------------------------+---------------------------------------+
   | Min-tls-version               | 1.2                                   |
   +-------------------------------+---------------------------------------+
   | SSH Mac-algorithm             | hmac-sha2-512-etm@openssh.com         |
   |                               | hmac-sha2-256-etm@openssh.com         |
   |                               | umac-128-etm@openssh.com              |
   |                               | hmac-sha2-256 hmac-sha2-512           |
   +-------------------------------+---------------------------------------+
   | SSH Key-exchange-algorithm    | curve25519-sha256                     |
   ```

```
|                                | curve25519-sha256@libssh.org         |
|                                | diffie-hellman-group-exchange-sha256 |
|                                | diffie-hellman-group16-sha512        |
|                                | diffie-hellman-group18-sha512        |
|                                | diffie-hellman-group14-sha256        |
+--------------------------------+--------------------------------------+
| SSH Cipher                     | non-cbc                              |
+--------------------------------+--------------------------------------+
| Telnet                         | Disabled                             |
+--------------------------------+--------------------------------------+
| Max-password-age               | 90                                   |
+--------------------------------+--------------------------------------+
| Force-default-password-change | Disabled                             |
+--------------------------------+--------------------------------------+

Common security settings
--- Time Elapsed: 127.976ms ---
```

b. View Global security settings with service-level overrides configured.

```
$ efa inventory device secure settings show

+--------------------------------+--------------------------------------+
|             Name               |                Value                 |
+--------------------------------+--------------------------------------+
| Min-tls-version                | 1.2                                  |
+--------------------------------+--------------------------------------+
| Min-service-tls-version        | HTTPS_CLIENTv1.3                     |
| [Applicable to OS ONE]         | LDAP_CLIENTv1.3                      |
+--------------------------------+--------------------------------------+
| SSH Mac-algorithm              | hmac-sha2-512-etm@openssh.com        |
|                                | hmac-sha2-256-etm@openssh.com        |
|                                | hmac-sha2-256 hmac-sha2-512          |
+--------------------------------+--------------------------------------+
| Telnet                         | Disabled                             |
+--------------------------------+--------------------------------------+
| Max-password-age               | 90                                   |
+--------------------------------+--------------------------------------+
| Force-default-password-change | Disabled                             |
+--------------------------------+--------------------------------------+

Common security settings
--- Time Elapsed: 24.952ms ---
```

c. View Device-Specific Settings.

```
efa inventory device secure settings show --ip <DEVICE_IP>
```

Example:

```
$ efa inventory device secure settings show --ip 10.64.188.63

+--------------------------------+--------------------------------------+
|             Name               |                Value                 |
+--------------------------------+--------------------------------------+
| Min-tls-version                | 1.2                                  |
+--------------------------------+--------------------------------------+
| Min-service-tls-version        | HTTPS_CLIENTv1.3                     |
| [Applicable to OS ONE]         | LDAP_CLIENTv1.3                      |
+--------------------------------+--------------------------------------+
| SSH Mac-algorithm              | hmac-sha2-512-etm@openssh.com        |
|                                | hmac-sha2-256-etm@openssh.com        |
|                                | hmac-sha2-256 hmac-sha2-512          |
+--------------------------------+--------------------------------------+
| SSH Cipher                     |                                      |
+--------------------------------+--------------------------------------+
| Telnet                         | Disabled                             |
```

```
+-------------------------------+---------------------------------------+
| Max-password-age              | 90                                    |
+-------------------------------+---------------------------------------+
| Force-default-password-change | Enabled                               |
+-------------------------------+---------------------------------------+
| AppState                      | cfg-in-sync                           |
+-------------------------------+---------------------------------------+

Current security settings on device 10.64.188.63
--- Time Elapsed: 24.416ms ---
```

Global settings are applied during device registration and can be updated at any time

4. Reset Security Settings.

   Reset one or more security settings to their factory default values. Each flag resets an individual setting independently — specify only the flags for settings you want to reset.

   ```
   efa inventory device secure settings reset --min-tls-version
   ```

   Note
   Resetting the minimum TLS version restores it to TLS 1.2 and removes all per-service overrides on OS ONE devices. Verify that all connected clients support TLS 1.2 before performing a reset.

Copy Configuration

You can copy the configuration.

About This Task
Follow the procedure to copy the configuration.

Procedure

1. Copy the default to the startup.

   OS ONE does not support this.

2. Copy the default to running.

   This command will erase the configuration from the device, but it will retain the certificates. The device will reboot as part of this command. The TPVM configuration will be preserved. You must reconfigure the gRPC server and associate the certificate to restore connectivity. Note that XCO does not natively support this operation. To perform this action, use the **`execute-cli`** command.

   ```
   efa inventory device execute-cli -ip 10.x.x.x --command="copy default-config running-config"
   ```

Factory Reset or Write Erase

This operation will erase the device's configurations and certificates. XCO does not natively support this operation. Use the **`execute-cli`** command to perform this task.

Monitoring XCO Status

The Monitoring service provides REST API for status, multiaccess, backup, restore and supportsave 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 32.
- Use the **`efa status`** command to verify application status.

  For more information, see the **`efa status`** command in the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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 |  APP STATE  |
+---------------------------------+-------+-------------+
| 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    | cfg-in-sync |
+---------------------------------+-------+-------------+
| Static Prefix Independent       | No    |             |
| Convergence                     |       |             |
```

```
+--------------------------------+-------+-------------+
| Maximum Load Sharing Paths     |       |             |
+--------------------------------+-------+-------------+
| Maximum Ipv6 Prefix Length 64  |       |             |
+--------------------------------+-------+-------------+
| Urpf                           |       |             |
+--------------------------------+-------+-------------+
| Ip Option Disable              | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Ip Option Disable Cpu          | No    |             |
+--------------------------------+-------+-------------+
| Ipv6 Option Disable            | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Peer Group Ipv6 Prefix Over    | Yes   | cfg-in-sync |
| Ipv4 Peer                      |       |             |
+--------------------------------+-------+-------------+
```

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 774.

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 125.

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: 20
- 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.

Note
The Database Restore operation does not support restoring a backup taken from a newer software version onto a system running an older version. The backup version must be equal to or older than the currently installed version on the target system. Attempting to restore a newer backup onto an older system will cause the restore to fail and leave the system in a non-operational state.

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 99.

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, 4.1.0*.

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, 4.1.0*.

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.

The `efa-monitor` service is available only on multi-node deployments. The logs contain:

- Galera and k3s recovery logs
- RabbitMQ recovery logs

The `efa-monitor` service checks status of clusters (galera, k3s or rabbitmq) every minute. If it detects any inconsistencies in the cluster, it will recover. The logs help in understanding the state of the cluster.

Use the **`sudo systemctl status/stop/start efamonitor`** to manage the monitoring service.

Before you manually stop any service for debugging purposes, ensure to stop the `efa-monitor` service. Otherwise, the `efa-monitor` service will automatically attempt to recover.

Enable or Disable Database Log Traces

You can use the **`efa fabric debug set`** command to set the debug mode for debugging purposes for the fabric service.

For information about commands and supported parameters to configure database log traces, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

The following table describes the database log trace settings in XCO 3.5.0 and XCO 3.6.0 and later:

Default Settings	XCO 3.5.0	XCO 3.6.0 and Later
Enable database log traces by default.	Default setting enables DB log traces whenever the fabric service starts or restarts.	Allows setting DB log levels and general-purpose logging levels.
Provide a CLI for enabling/ disabling database log traces and setting debug levels for general-purpose logging.		

The following table describes the log level and the CLI output information:

User Configuration	Log Level	Information Reflected in CLI Output
Users can set the debug mode for the fabric service.	1. **debug:** Used for setting verbose level logging. 2. **info:** Used for setting general level logging. 3. **debugdb:** Used to enable SQL traces to the log file. 4. **nodebugdb:** Used to disable SQL traces.	• User history data. • Flag levels: debug, info, debugdb, nodebugdb.

Back up and Restore the XCO System

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

About This Task

Follow this procedure to back up and restore the XCO system.

Procedure

1. To back up the system, complete the following steps:
   a. (If remote option is not specified) To back up the system, run the following command:

```
$ efa system backup
Generating backup of EFA...
Backup Location: /apps/efa_logs/backup/EFA-2023-09-10T13-21-46.413.tar
--- Time Elapsed: 10.401999384s ---
```

   b. (If remote option is specified) To back up the system and copy to remote server after configuring remote server settings, run the following command:.

```
$ efa system backup --remote

Generating backup of EFA...

Backup Location on local server: /var/log/efa/backup/
EFA-3.4.0-1102-2023-08-11T17-56-46.940.tar
Backup Location on Remote Server: user@10.37.34.151:/home/user/
EFA-3.4.0-1102-2023-08-11T17-56-46.940.tar

--- Time Elapsed: 18.462460717s ---
```

2. To restore the system, complete the following steps:
   a. Run the **`efa system restore`** command.

Note

As a best practice, always restore on a fresh XCO install with no multi-access or management sub-interface config.

```
$ efa system restore

EFA-2023.01.12-11.19.52.tar (Version:2.3.2-GA, Generated by: User)
EFA-2023-01-12T04.59.09.tar (Version:2.4.0-7171, Generated by: User)
EFA-2023-01-12T13.50.00.tar (Version:2.4.0-121, Generated by: User)
EFA-2023-01-12T13.50.51.tar (Version:2.4.0-121, Generated by: System)
```

```
EFA-2023-01-12T16.11.47.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-2023.01.12-11.19.52.tar
Performing EFA restore using EFA-2023.01.12-11.19.52.tar
Generating backup before initiating restore
BACKUP_TAR: /apps/efa_logs/backup/EFA-2023.01.12-11.19.52.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 --force
```

The command installs certificates on SLX devices, and can be applied to one or more devices.

Note
For Extreme OS ONE, JWT and gRPC certificates are installed, whereas HTTPS and OAuth2 certificates are installed on SLX 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.

Recover XCO and End-Device Configuration

You can retrieve both the XCO configuration and the end-device configuration.

About This Task

Follow this procedure to retrieve both the XCO configuration and the end-device configuration. Complete the backup procedure first, and then complete the restore procedure.

Procedure

1. Capture the Backup.

   See *Periodic Backups and Configuration* under XCO System Backup and Restoration on page 96.

2. Restore and Replay.
   a. Freshly deploy XCO in HA mode, using the same build, on the TPVMs that reside on both spines.
   b. Copy the backup file from the remote location to the `/apps/efa_logs/backup/` directory.

      Ensure that the copied file uses the naming format `EFA-4.1.0-92-2026-08-03T13-51-43.593.tar`. Ensure that the backup filename maintains the original naming convention
   c. Restore the XCO backup. Run the **`efa system restore`** command and select the backup file that you created during the capture procedure.
   d. Monitor the restore process. After the services come up, install the device certificates.

      ```
      efa certificate device install --fabric <fabric-name> --force
      ```
   e. Wait approximately 300 seconds for services to stabilize.
   f. Run the show commands and confirm that all entities (fabric, PO, VRF, EPG, PG, and peer) are in a refreshed state with health shown as BLACK.
   g. Run the c**`onfig-backup history`** command and record the UUID and SSID for each device.
   h. Start config-replay for all left-hand leaf and BL devices at once, using the `--startup-config` token.

      ```
      efa inventory config-replay execute --ip <standby_switch_ip> --uuid
      <uuid_of_slx_backup> --startup-config
      ```
   i. Monitor the device status. After the devices come up, confirm that they enter Maintenance Mode (MM) and that DRC is triggered.
   j. Monitor DRC progress and verify that no drift is detected, because the replay should have pushed all configurations. Confirm that the devices exit MM after a successful DRC run.
   k. Start config-replay for all right-hand leaf and BL devices at once, using the --startup-config token.

      ```
      efa inventory config-replay execute --ip <standby_switch_ip> --uuid
      <uuid_of_slx_backup> --startup-config
      ```
   l. Monitor the device status. After the devices come up, confirm that they enter MM and that DRC is triggered.
   m. Monitor DRC progress and verify that no drift is detected. Confirm that the devices exit MM after a successful DRC run.
   n. Start config-replay for the standby device, using the --startup-config token so that no failover is triggered.

      ```
      efa inventory config-replay execute --ip <standby_switch_ip> --uuid
      <uuid_of_slx_backup> --startup-config
      ```
   o. Monitor the device status. After the device comes up, confirm that it enters MM and that DRC is triggered.
   p. Monitor DRC progress and verify that no drift is detected. Confirm that the device exits MM after a successful DRC run.
   q. After DRC completes for the standby node, reload the active-node TPVM to trigger a failover.

r. After all services come up following the failover, start config-replay for the current standby node (the previous active node), using the --startup-config token so that no failover is triggered.

```
efa inventory config-replay execute --ip <standby_switch_ip> --uuid
<uuid_of_slx_backup> --startup-config
```

s. Monitor the device status. After the device comes up, confirm that it enters MM and that DRC is triggered.

t. Monitor DRC progress and verify that no drift is detected. Confirm that the device exits MM after a successful DRC run.

u. Run the show commands and confirm that all components are back in sync and that XCO health and fabric health return to GREEN.

Supportsave Enhancement

XCO has introduced delimiters to encapsulate command outputs when collecting CLI results. This upgrade enables the capture of additional command outputs, making it easier to retrieve data directly from commands instead of relying on database outputs.

Key improvements include:

- Adding delimiters at the start and end of command outputs for clearer encapsulation
- Collecting additional command outputs to provide more comprehensive data
- Enabling direct data retrieval from commands, reducing reliance on database outputs
- Timestamps have been included for all command output in the SupportSave records
- Adding timestamps for all the command outputs

Example output:

The following are the example output post the supportsave enhancements:

```
*******************************************************
CLI START: efa status
Thu Feb 06 07:59:16 UTC 2025
*******************************************************
EFA application status:
+-----------+--------+--------+--------------+
| Node Name | Role   | Status | IP           |
+-----------+--------+--------+--------------+
| xco-vm-12 | active | up     | 10.64.193.79 |
+-----------+--------+--------+--------------+
--- Time Elapsed: 1.202930986s ---
CLI END: efa status

*******************************************************
CLI START: efa inventory device list
Thu Feb 06 07:59:17 UTC 2025
*******************************************************
Inventory listed device in EFA:
+----------------+-----------+-------+--------------+----------+------------+------
+-------------+-------------------------+-------------------------+----------------
+-------------+-----------------+-----------------------------+-------------------
+-----------------+--------+----------+
|   IP Address   | Host Name | Model | Chassis Name
| Firmware |    ASN     | Role |   Fabric    |    RegistrationTime
|    LastDiscoveryTime    | Cert/Key Saved | Admin State | Maintenance Mode
```

```
| Maintenance Mode on Reboot | Syslog Registered | SNMP Registered |  Type  | Location |
+----------------+-----------+-------+-------------+----------+------------+------
+-------------+-------------------------+-------------------------+-----------------
+-------------+------------------+--------------------------+------------------
+-----------------+--------+----------+
| 192.168.246.12 | b182_L02  | 3200  | 8720-32C     |
20.6.2a  | 4200000000 | Leaf | SDI3-FABRIC | 2024-09-02
10:57:47 UTC | 2024-09-03 15:56:05 UTC | Y              | up          | Disable
| Enable                      | Y                | Y              | FABRIC | default  |
+----------------+-----------+-------+-------------+----------+------------+------
+-------------+-------------------------+-------------------------+-----------------
+-------------+------------------+--------------------------+------------------
+-----------------+--------+----------+
| 192.168.246.11 | b182_L01  | 3200  | 8720-32C     |
20.6.2a  | 4200000000 | Leaf | SDI3-FABRIC | 2024-09-02
10:57:47 UTC | 2024-09-03 15:52:15 UTC | Y              | up          | Disable
| Enable                      | Y                | Y              | FABRIC | default  |
+----------------+-----------+-------+-------------+----------+------------+------
+-------------+-------------------------+-------------------------+-----------------
+-------------+------------------+--------------------------+------------------
+-----------------+--------+----------+
Device Details
--- Time Elapsed: 17.832694ms ---
CLI END: efa inventory device list
```

The following example shows the additional command output to be collected. These commands output will be added to the `advance_cli_data.txt`:

```
efa fabric show --detail
efa fabric execution show
efa inventory device list
efa inventory drift-reconcile history
efa inventory execution show
efa inventory firmware-host list
efa inventory rma history
efa inventory config-backup history
efa inventory config-replay history
efa inventory admin-state history
efa tenant show
efa tenant po show
efa tenant epg show --detail
efa tenant vrf show --detail
efa tenant service bgp peer show
efa tenant service bgp peer-group show
efa tenant execution show
efa tenant service mirror
efa notification subscribers list
efa snmp subscriber list
efa snmp execution show
efa scvmm list
efa vcenter list
efa openstack network show
efa openstack subnet show
efa openstack network-interface show
efa openstack router-interface show
efa openstack router-route show
efa openstack router show
efa openstack execution show
efa auth client show
efa auth rolemapping show
efa auth ldapconfig show
efa auth settings token show
efa auth execution show
efa system backup-list
```

```
efa system execution show
efa system settings show
efa system feature show
efa system restore-history show
efa status
efa system alert show
efa system alarm show
efa health detail show
efa health inventory show
efa mgmt subinterface show
efa mgmt route show ipv4
efa mgmt route show ipv6
efa policy prefix-list list --type ipv4
efa policy prefix-list list --type ipv6
efa policy route-map list
efa policy community-list list
efa policy extcommunity-list list
efa policy large-community-list list
efa policy qos map list
efa policy qos service-policy-map list
efa policy qos profile list

efa inventory device firmware-download history --type fabric
efa inventory device tpvm list
efa inventory kvstore list
efa inventory debug show-fwdl-Inprogress-devices
efa inventory debug device-adapter-status --fabric-all
```

The following example shows the option to exclude advanced or additional CLI output is available when collecting support saves:

```
efa system supportsave --disable-advance-cli-capture
(efa:user)user@xco-vm-12:~$ efa system supportsave -h
Generate supportsave for EFA

Usage:
  efa system supportsave [flags]

Flags:
      --device-ip string              Comma separated list of device IP Address to
collect supportsave from
      --disable-advance-cli-capture   Disable the collection of advance CLI output.
      --fabric-all                    supportsave for all devices
      --fabric-name string            supportsave of all devices part of this fabric
  --- Time Elapsed: 389.268µs ---
(efa:user)user@xco-vm-12:~$ efa system supportsave --disable-advance-cli-capture
SupportSave File Location: /var/log/efa/efa_2025-01-31T07-02-34.657.logs.zip
--- Time Elapsed: 10.477560239s ---
(efa:user)user@xco-vm-12:~$
```

Data Collection and Troubleshooting Enhancements

Additional Debugging Information During Execute CLI

XCO offers the ability to run device CLI commands through its own CLI option. However, commands may fail due to device or connectivity issues, often resulting in a generic error message stating "the device is not reachable," even when it is.

To enhance this, XCO provides additional debugging information to help users identify the underlying issues. Command execution failures may originate from SSH errors, REST call failures, or HTTP server problems. With this enhancement, users will receive specific feedback

to address the issues effectively. For instance, if a command fails due to an HTTP server issue on the device, the output will provide the following details:

```
(efa:extreme)extreme@ok1c1tpvm1:~$ efa inventory device execute-cli --ip 10.249.42.199 --
command "show clock"

Execute CLI[failed]
------------------------------------------------------------------------------------------
------------------------------------------------------------------------
IP Address     Host Name     Fabric     Command     Status     Reason     Output
------------------------------------------------------------------------------------------
------------------------------------------------------------------------
10.249.42.199     OK2C1T02     OK1C1T01-OK2C1T02     show clock     Failed     Device
10.249.42.199 Failed to run the command as rest execution failed.
------------------------------------------------------------------------------------------
------------------------------------------------------------------------
Execute CLI Details
PS: check http server status on the device.
— Time Elapsed: 2m2.092072921s —



(efa:extreme)extreme@ok1c1tpvm1:~$ efa inventory device execute-cli --ip 10.249.42.199 --
command "show clock"
Execute CLI[failed]
------------------------------------------------------------------------------------------
------------------------------------------------------------------------
IP Address      Host Name  Fabric Command     Status      Reason   Output
------------------------------------------------------------------------------------------
------------------------------------------------------------------------
10.249.42.199  OK2C1T02    OK1C1T01-OK2C1T02  show clock  Failed   Device 10.249.42.199
Failed to run the command as rest execution failed.

------------------------------------------------------------------------------------------
------------------------------------------------------------------------
Execute CLI Details
PS: check http server status on the device.
— Time Elapsed: 2m2.092072921s —
```

Additional Device Side Commands Collection During XCO Supportsave

When the supportsave for XCO is collected, additional command outputs will be gathered for each device and stored in individual files named `<device_ipaddr>_info.txt`. Each file will include outputs from multiple commands, separated by delimiters for clarity. This collection is disabled by default and can be enabled using the `--collect-device-information` option.

The collected data will be saved under a new folder named `device_data/` within the support save directory structure.

Currently, XCO collects the device supportsave and transfers it to a remote location when options such as `--device-ip`, `--fabric-all`, or `--fabric-name` are specified. Starting with XCO release 4.0.2, if the `--collect-device-information` option is used, XCO will only collect additional command outputs from the specified devices (as determined by the

`--device-ip`, `--fabric-name`, or `--fabric-all` options) without triggering the standard device support save.

Note
The `--collect-device-information` option must be used in combination with one of the following options: `--device-ip`, `--fabric-name`, or `--fabric-all`.

The following table provides the further information on the CLI commands and their supportsave collection behavior:

XCO support save CLI Command	Support save collection behavior
`efa system supportsave` along with any of the following option only: 1. --device-ip 2. --fabric-all 3. --fabric-name	Collects the device support save and copies to remote server if configured.
`efa system supportsave` along with any of the following option along with **`--collect-device-information`**: 1. --device-ip 2. --fabric-all 3. --fabric-name	Collects the additional commands output from device which is mentioned in the **List of additional device commands**. The additional command output will be with the XCO support save. Even though the filter is provided, it will not collect the device support save.
`efa system supportsave` with no filter option	Collects the XCO support save only

The following is a sample command output:

```
Path: device_data/
File Name: 10.32.43.233_info.log
*****************************************************************
CLI START: show bgp ip flowspec neighbors
Wed May 14 08:56:01 UTC 2025
*****************************************************************
Local Port    Dead Interval  Remaining Life  Remote Port ID   Remote Port Descr   Chassis
ID       Tx      Rx      System Name
Eth 0/1       120            99              ethernet 0/5     Ethernet 0/5
f46e.959f.1708   49215   174     TOS164-12
Eth 0/10      120            117             ethernet 0/20    Ethernet 0/20
f46e.959f.1708   49220   176     TOS164-12
Eth 0/11      120            90              Ethernet 0/11    Eth 0/11
78a6.e145.5c14   51296   2609    SLX-164-44
Eth 0/12      120            90              Ethernet 0/12    Eth 0/12
78a6.e145.5c14   51296   2609    SLX-164-44
Eth 0/13      120            90              Ethernet 0/13    Eth 0/13
78a6.e145.5c14   51296   2609    SLX-164-44
Eth 0/14      120            90              Ethernet 0/14    Eth 0/14
78a6.e145.5c14   51296   2609    SLX-164-44
Total no. of Records: 6
CLI END: show lldp neighbors detail
*****************************************************************
CLI START: show interface stats brief
Wed May 14 08:56:01 UTC 2025
*****************************************************************
```

```
                Packets                    Error             Discards          CRC
Interface    rx          tx          rx          tx      rx          tx      rx
=========    ==========  ==========  ==========  ======  ==========  ======  =====
Lo 1         0           0           0           0       0           0       0
Lo 2         0           0           0           0       0           0       0
Eth 0/1      0           0           0           0       0           0       0
Eth 0/2      0           0           0           0       0           0       0
Eth 0/3      0           93414       0           0       0           0       0
Eth 0/4      0           0           0           0       0           0       0
Eth 0/5      0           0           0           0       0           0       0
Eth 0/6      0           0           0           0       0           0       0
Eth 0/7      0           0           0           0       0           0       0
Eth 0/8      0           0           0           0       0           0       0
Eth 0/9      0           0           0           0       0           0       0
Eth 0/10     0           0           0           0       0           0       0
```

- **List of additional device commands**
  - show bgp evpn summary
  - show ip bgp summary
  - show cluster
  - show tunnel brief
  - show overlay-gateway vni-mappings
  - show overlay-gateway
  - show lldp neighbors
  - show bfd neighbors
  - show bfd
  - show ip interface brief
  - show vlan brief
  - show bridge-domain
  - show mac-address-table
  - show ip arp
  - show port-channel
  - show interface stats brief
  - show system maintenance
- **CLI command option**

  The additional filter option is included to enhance data collection from the device.

```
(efa:user)user@xco-vm-12:~$ efa system supportsave
Generate supportsave for EFA
 Usage:
  efa system supportsave [flags]

Flags:
    --device-ip string              Comma separated list of device IP Address to
collect supportsave from
    --disable-advance-cli-capture   Disable the collection of advance CLI output.
    --collect-device-information    will collect the additional device side CLI
information during support save.
    --fabric-all                    supportsave for all devices
    --fabric-name string            supportsave of all devices part of this fabric

--- Time Elapsed: 389.268µs ---
```

Installation Type and System Uptime in XCO Supportsave

Installation type (OVA or Server) and system uptime are included in the XCO supportsave file.

The following is a sample output:

```
Path: misc/
File: system_info_<ip>.log
Installation:
<OVA>/<Server>
Uptime:
07:09:58 up 14:15,  3 users,  load average: 0.25, 0.20, 0.18
```

Supportsave Collection Duration in Large-Scale Deployments

Supportsave collection may take longer in large-scale configuration environments, especially when services such as inventory, fabric, and tenant are busy processing configuration updates in the background.

Factors Affecting Supportsave Duration

The following factors can increase the time required for supportsave collection:

- **Large-scale configurations**: Environments with a high number of VRFs, EPGs, and BGP peer groups require more time for data collection. For example, configurations with 200 VRFs, 1000 EPGs, and 400 BGP peer groups may take approximately 7 minutes even when the system is idle.
- **Background service activity**: If supportsave is executed while services such as Inventory, Fabric, or Tenant are actively processing events, including Drift-Reconcile (DRC) operations or configuration updates, the collection duration can increase significantly due to CPU and database resource contention.
- **Advance CLI capture**: The advance CLI capture feature is enabled by default and collects additional diagnostic information from services, which adds to the overall supportsave duration.

Reducing Supportsave Duration

To reduce the supportsave collection time:

- Disable the advance CLI capture by using the **`--disable-advance-cli-capture`** command when time-sensitive collection is needed and full service-level diagnostics are not required:

  ```
  efa system supportsave --disable-advance-cli-capture
  ```

  

  Note
  Disabling advance CLI capture excludes certain service-level diagnostic data from the supportsave bundle. Use this option only when faster collection is required and the additional diagnostic data is not needed.

- Run supportsave when the system is idle, with no active DRC, firmware download, or large-scale configuration operations in progress, for optimal collection times.
- If supportsave must be collected during active operations, expect longer collection times proportional to the system load.

Passwordless SSH or SCP Support for Secure and Efficient Backup and Log File Operations

XCO introduces SSH key-based authentication as an alternative to traditional SCP credential handling, significantly enhancing security by addressing password-related vulnerabilities. The SSH key-based passwordless authentication between XCO and remote servers ensures secure interaction with remote servers without repeated password entry.

Key Benefits

- Safeguards sensitive data during file transfers and communications
- Preferred for its direct integration with the SSH protocol, ease of use, and simplicity

Supported Features

- Generate Key Pair: XCO prompts users to generate a new SSH key pair, specifying options like RSA algorithm, key size, and passphrase.
- Store Private Key: XCO securely stores the SSH key pair in the database.
- Copy Public Key to Server: XCO provides CLI support to display the public key for manual copying to the remote server.
- Set Up Remote Server Connection: XCO saves connection details and associates the SSH key pair with the connection.
- Upload SupportSave and Logging Files: XCO initiates an SSH connection using the stored key pair for authentication and file copy operations.
- Authentication Success: XCO gains access to the server if the keys match, allowing upload and logging operations to proceed.

Error Scenario

1. Updating setting with username other than that of the username used while copying public key. SSH public key copied for username "user" and tried to access using username "root".

```
(efa:sbr)sbr@efa1:~/files$ efa system settings update --remote-server-ip 10.32.85.29
--remote-server-username root --remote-server-directory /home/user/efa_test/ --
passwordless Keybased --ssh-key-id rsa_user
Error :  SSH key-based authentication failed: Error creating session on remote: ssh:
handshake failed: ssh: unable to authenticate, attempted methods [none publickey], no
supported methods remain
```

2. Updating settings with wrong ssh-keyid

```
efa:sbr)sbr@efa1:~/files$ efa system settings update --remote-server-ip 10.32.85.29
--remote-server-username root --remote-server-directory /home/user/efa_test/ --ssh-key-
id rsa_root
Error: Fetching SSH KeyID failed: SSH key ID: rsa_root does not exist
```

3. If wrong public key copied to remote server authorized_keys file

```
efa:sbr)sbr@efa1:~/files$
efa system settings update --remote-server-ip 10.32.85.29
--remote-server-username root --remote-server-directory /home/user/efa_test/
--ssh-key-id id_rsa

Error :  SSH key-based authentication failed: Error creating session on remote: ssh:
handshake failed: ssh: unable to authenticate, attempted methods [none publickey], no
supported methods remain
```

Configure SSH Key

About This Task

Use the **`efa system ssh-key`** command to generate, show, update, and delete SSH keys.

Procedure

1. Using CLI: Complete the following task.
   a. To generate the SSH key, run the following command:

```
(efa:user)user@pkumarpatra32:~$ efa system ssh-key generate --algorithm-type rsa --
ssh-key-id my_rsa_key --key-size 2048
Enter passphrase (leave empty for no passphrase):
Public Key: ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABAQCwouqscy4CkkWVAoLEw6F9xFnOtOV79bxMU6sAlD9KIev1KPv8amni
e80qYdzb9a9DxIx5tTbBMfBiHbvMjMNoJPQ3ixw50BgoFLI7LqDPTTUu064xobH60Ctpwe7ispSH6XbwiTlx
Uj9vWWwWG9NlxNHncXOX6ARXWSbj5J+L0p35RVbck5LQV9FIcRz+hUy8muXiPFEyOQyZN7hWzY5uCPKlTMYI
ZJoPnKyxJR+EW4A1Hin3V60ray9Lfu0MkP9w9OETnyRzc9ctyqpuV+uPM8cFhfYOBw+ZSLx36ScSbX+ddMV6
hnVtfk7dctlsgvo0JxFpwD2jmMDKNONiLx1x
--- Time Elapsed: 1.610456842s ---
```

   A unique SSH key pair, consisting of a private key and a public key, will be generated based on the chosen algorithm and key size. The pair will be securely stored in the database, with the private key and passphrase encrypted. To maintain efficiency, a maximum limit of 20 SSH keys per user is enforced.

Note

If a key pair is already associated with a user, it can be regenerated. However, it is the user's responsibility to update the new public key on their remote server.

Field Name	Data Type	Required	Description
ssh-key-id	String	true	A unique identifier for the SSH key pair. Example: "my_rsa_key"
algorithm-type	String	true	Specifies the algorithm used to generate the SSH key pair. Supported algorithms: rsa, ecdsa, ed25519. Example: "rsa"
key-size	int	false	The size of the key in bits when generating the SSH key. Supported Key Sizes: • RSA: 1024, 2048, 4096, 8192 bits (default: 2048) • ECDSA: 256, 384, 521 bits (default: 256) • Ed25519: Fixed at 256 bits (default: 256 • Example: 2048
passphrase	String	false	An optional string to secure the private key. If omitted, the private key is stored unencrypted. Example: "securepassword123"

The following are some Error Scenarios while generating the SSH key::

i. Missing algorithm_type: The client sends a request without the algorithm_type field.

```
(efa:user)user@pkumarpatra32:~$ efa system ssh-key generate --ssh-key-id
my_rsa_key --key-size 2048
Error :  algorithm type is mandatory
```

ii. Invalid algorithm_type: The client sends a request with an invalid algorithm_type value (i.e., a value that is not rsa, ecdsa, or ed25519).

```
(efa:user)user@pkumarpatra32:~$ efa system ssh-key generate  --ssh-key-id
my_rsa_key --algorithm-type rs --key-size 2048
Enter passphrase (leave empty for no passphrase):
Error :  Algorithm type 'rs' is not supported. Supported algorithms: rsa, ecdsa,
ed25519.
```

iii. Invalid key_size for the selected algorithm: The client provides an invalid key_size value that is not supported by the selected algorithm_type.

```
(efa:user)user@pkumarpatra32:~$ efa system ssh-key generate  --ssh-key-id
my_rsa_key --algorithm-type rsa --key-size 123
Enter passphrase (leave empty for no passphrase):
Error :  invalid key size for RSA. Supported sizes: 1024, 2048, 4096, 8192 bits
```

iv. Maximum SSH Keys Limit Exceeded: If the user has already created maximum keys, an error message will be returned when they attempt to create a new key.

```
(efa:user)user@pkumarpatra32:~$ efa system ssh-key generate  --ssh-key-id
my_rsa_key1 --algorithm-type rsa

Error: maximum number of SSH keys reached: 20
```

v. Internal Server Error (Unexpected System Failure): There is an unexpected failure in the system, such as a database issue or SSH key generation failure.

```
(efa:sbr)sbr@efa1:~$efa system ssh-key generate --algorithm-type rsa -ssh-keyid
my_rsa_key1 --key-size 2045
Error: Internal server error. Please try again later.
```

b. To display the SSH key, run the following command.

Displaying the public key allows users to easily retrieve the public key generated earlier.

```
(efa:sbr)sbr@efa1:~$ efa system ssh-key show
+------------+----------------+-------------+
| SSH Key ID | Algorithm type | SSH Key Size|
+------------+----------------+-------------+
| my_rsa_key |    rsa         |  2048       |
+------------+----------------+-------------+
|my_ecdsa_key|   ecdsa        |   256       |
+------------+----------------+-------------+
|my_rsa_key1 |    rsa         |  1024       |
+------------+----------------+-------------+
--- Time Elapsed: 95.469328ms ---
```

To verify the SSH key details based on the specific ssh_keyid, run the following command:

```
(efa:sbr)sbr@efa1:~$ efa system ssh-key show --ssh-key-id ecdsa1
Public Key: ecdsa-sha2-nistp256
AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBDp0YT1ZlnrFYD1mL+fTeAWoXQq7/
FrMqqvjFjqFHmY0j1k8ra+32gxstssXucaoXQ53NmoXQEbJNw0qxcPBi+M=
--- Time Elapsed: 28.446504ms ---
```

Error Scenarios

- If no SSH key exists for a specific ssh_keyid, the response will return an error message.

```
$ efa system ssh-key show --ssh-key-id my_rsa_key1
Error :  SSH key ID: my_rsa_key1 does not exist
```

c. To export the public key as .pub file which can be exported to a remote server, run the following command:

```
efa system ssh-key export --ssh-key-id my_rsa_key
File Exported Successfully
File: my_rsa_key.pub
Location: /home/user/my_rsa_key.pub
--- Time Elapsed: 32.126921ms ---
```

The public key file will be displayed as a response.

d. To delete the SSH key, urn the following command:

```
(efa:user)user@pkumarpatra32:~$ efa system ssh-key delete --ssh-key-id my_rsa_key
SSH Key pair my_rsa_key deleted successfully
--- Time Elapsed: 37.141292ms ---
```

The following are some error scenarios while deleting the SSH key:

i. If the key pair is in use and cannot be deleted, the response will return an error message:

```
(efa:sbr)sbr@efa1:~$ efa system ssh-key delete --ssh-key-id my_rsa_key

Error: SSH key pair cannot be deleted because it is in use
```

ii. If the specified ssh_keyid does not exist in the system:

```
$ efa system ssh-key delete --ssh-key-id my_rsa_key
Error :  SSH key ID: my_rsa_key does not exist
```

2. Using REST API: Complete the following task.

a. To generate the SSH key, use the following API:

```
curl --location --request PUT 'http://gosystem-service:80/v1/system/ssh-key'
--header 'Content-Type: application/json' \
--data-raw '{
      "ssh-key-id": "my_rsa_key",
      "algorithm-type": "rsa",
      "passphrase": "pwdabc",
      "key-size": 2048
}'
```

Example Response Body:

If the SSH key pair is generated successfully, the system responds with an HTTP 201 status code and includes the corresponding public key

```
{
"public_key": "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABgQDSDb0QTKr+JBVuKXhHhj0ANck3VMmulaitTniROJtbCwyzwY9ysQko
yAix8HMtpLT0rvMXNIJZ/F2v/+7asxB5uJWQHR0K57fAja82cMAZS0wU/
Nj3crPc3ise08NK78EkXBuKEpwt6PMVgZMYRnFLMd2NZTvWv1zet1kMIxtA/
Fn3QWejlfKa0tovOCk5TxtyYc5uIipHkkqAKM/
YlTlNyQlLa9Wnbz56HmMJ7PZJq7tIK6HE8jRSo+mPLXTolwbtlQG+AADwYDyrzExOVUZrqHWGvkkdZNypv2/
/hN9qPEecT1RbNati1OzIvw6aEQiYjCwAT4GiRg2stADqekOolV5rQCQ4sAGmBwgsdHc58BNP/
94ZSWnNkLxm3NIyHKw43c9OqfOfohsescmFE+EVi8s/RQsIYiUOZf8rSLuzYVXoG2IQV/
MYs7lTanV7YS48jkR5GRR+QuzPIFhLo9UfyUL3LeJHUB6HtoidcY0337g2LhgBp8GTWbSbW5r/w2k="
}
```

b. To display all the available SSH keys, use the following API:.

```
curl --location --request GET 'http://gosystem-service:80/v1/system/ssh-key'
```

To display the ssh-key based on the ssh_keyid, use the following REST API:

```
curl --location --request GET 'http://gosystem-service:80/v1/system/ssh-key?ssh-key-id=my_rsa_key'
```

Query Parameter: ssh_keyid (string, required): The ID of the SSH key pair for which the user wants to fetch the corresponding public key.

Example Response Body:

If the public key is found for a specific ssh_keyid, the response will look like this:

```
{
          "ssh-key-id": "my_rsa_key",
          "algorithm-type": "rsa",
          "key-size": 2048,
"public-key":"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABgQDSDb0QTKr+JBVuKXhHhj0ANck3VMmulaitTniROJtbCwyzwY9ysQko
yAix8HMtpLT0rvMXNIJZ/F2v/+7asxB5uJWQHR0K57fAja82cMAZS0wU/
Nj3crPc3ise08NK78EkXBuKEpwt6PMVgZMYRnFLMd2NZTvWv1zet1kMIxtA/
Fn3QWejlfKa0tovOCk5TxtyYc5uIipHkkqAKM/
YlTlNyQlLa9Wnbz56HmMJ7PZJq7tIK6HE8jRSo+mPLXTolwbtlQG+AADwYDyrzExOVUZrqHWGvkkdZNypv2/
/hN9qPEecT1RbNati1OzIvw6aEQiYjCwAT4GiRg2stADqekOolV5rQCQ4sAGmBwgsdHc58BNP/
94ZSWnNkLxm3NIyHKw43c9OqfOfohsescmFE+EVi8s/RQsIYiUOZf8rSLuzYVXoG2IQV/
MYs7lTanV7YS48jkR5GRR+QuzPIFhLo9UfyUL3LeJHUB6HtoidcY0337g2LhgBp8GTWbSbW5r/w2k="
}
```

You can copy the SSH public key to the `~/.ssh/authorized_keys` file on remote servers.

c. To delete the SSH key, use the following API:

Query Parameter: ssh_keyid (required): The unique identifier for the SSH key pair that the user wishes to delete.

```
curl --location --request DELETE 'http://gosystem-service:80/v1/system/ssh-key?ssh-key-id=my_rsa_key
```

Example Response Body:

If the record is successfully deleted the API will return http accepted (202).

Associate SSH Key Pair for Secure Operations

About This Task

Associating an SSH key pair with system operations (for example, periodic backups and support save retrieval) enables secure and passwordless authentication for remote server interactions, streamlining backup and support save retrieval.

By integrating SSH key pairs, system operations like backups and support saves can be securely authenticated and transferred to remote servers without manual password entry.

Procedure

1. Using CLI: To associate SSH key Id with remote settings, run the following command:

```
efa system settings update --remote-server-password password --remote-server-ip
10.32.85.25 --remote-server-username root --remote-server-directory /home/sbr/sspath --
ssh-key-id my_rsa_key
```

Note

Password is still required to collect the device support save. Once the key is associated, the system uses the corresponding private key to authenticate with the remote server, enabling secure operations such as scheduled backups and file uploads.

```
SSHKeyId: KeyPair name should be passed as the value
      Example: my_rsa_key
```

a. Once the association of SSH key id is completed, run the following show command:

```
(efa:extreme)extreme@tpvm-122:~$ efa system settings show
+-------------------------------+-----------------------------+
|             SETTING           |     VALUE                   |
+-------------------------------+-----------------------------+
| Max Backup File Limit         |        5                    |
+-------------------------------+-----------------------------+
| Max Supportsave File Limit    |        5                    |
+-------------------------------+-----------------------------+
| Backup Schedule               |     0:* :*:*                |
+-------------------------------+-----------------------------+
| Remote Server Ip              |    10.32.85.25              |
+-------------------------------+-----------------------------+
| Remote Server Username        |     root                    |
+-------------------------------+-----------------------------+
| Remote Server Password        |    KCr+INRnovA=             |
+-------------------------------+-----------------------------+
| Remote Server Directory       |   /home/sbr/sspath          |
+-------------------------------+-----------------------------+
| Remote Transfer Protocol      |       scp                   |
+-------------------------------+-----------------------------+
| Periodic Device Config Backup |     Disabled                |
+-------------------------------+-----------------------------+
| Show User Creation On Startup |       true                  |
+-------------------------------+-----------------------------+
| Private Key                   |                             |
+-------------------------------+-----------------------------+
| Certificate Key               |                             |
+-------------------------------+-----------------------------+
| CA Public Key                 |                             |
+-------------------------------+-----------------------------+
| Passphrase                    |                             |
+-------------------------------+-----------------------------+
| SSHKeyID                      |      my_rsa_key             |
```

```
+------------------------------+--------------------------+
--- Time Elapsed: 51.469328ms ---
```

Note

- To use the available remote settings, reset them before updating.
- To switch from certificate-based authentication to key-based authentication, or vice versa, you must reset the settings first before attempting again.
- The certificate-based authentication and key-based authentication cannot be configured or supported simultaneously.

2. Using REST API: To associating SSH key Id with remote settings, use the following API:

```
curl --location --request POST 'http://gosystem-service:80/v1/system/settings' \
--header 'Content-Type: application/json' \
--data-raw '{
  "keyval": [
    {
      "value": "5",
      "key": "MaxBackupFiles"
    },
    {
      "value": "0 0 * *",
      "key": "BackupSchedule"
    },
    {
      "value": "5",
      "key": "MaxSsFiles"
    },
    {
      "value": "10.10.10.10 / 2000::1",
      "key": "RemoteServerIP"
    },
    {
      "value": "scp / ftp",
      "key": "RemoteTransferProtocol"
    },
    {
      "value": "username",
      "key": "RemoteServerUsername"
    },
    {
      "value": "password",
      "key": "RemoteServerPassword"
    },
    {
      "value": "/root/test",
      "key": "RemoteServerDirectory"
    },
    {
      "value": "Enabled",
      "key": "PeriodicDeviceConfigBackup"
    },
    {
      "value": "my_rsa_key",
      "key":  "SSHKeyID"
    },
  ]
}
'
```

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 inactive during this step. In a multi-node deployment, XCO stays active when the command is run on the standby node but becomes inactive if the command is run 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 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, complete the following steps:
   a. Change the directory.
      - On TPVM-based deployment
        ```
        cd /apps/bin/
        ```
      - On Server-based deployment
        ```
        cd /usr/local/bin
        ```
   b. Run the following command:
      ```
      sudo bash efa-change-vip <New-VIP>
      ```

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

Note
From XCO v4.1.0, the **`efa show-running-config`** command only shows interface attributes (MTU, Speed, FEC, admin-state, etc.) if explicitly configured by user.

- Device-default values are not shown.

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 <username> --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 <username> --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.

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`

  K3s service is available on single and multi-node deployments. The log contains journalctl log of k3s for the past one day.
- Keepalived service log files: `/var/log/keepalived`. The directory contains keepalived service logs. You can use keepalived service logs to debug failovers, double faults, and gateway connectivity. The log contains
  - Journalctl log of keepalived for the past one day
  - Keepalived service logs

  Keepalived service is only available in multi-node deployments. Keepalived has three states: MASTER, BACKUP and FAULT. The active node must have the status of MASTER, the standby will have BACKUP, and FAULT can be seen during a failover.

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.

It captures logs of micro-services in XCO, services deployed on the host, and a snapshot of the database. The logs from executing the Supportsave command can be found in the monitoring service logs located at `/apps/efa_logs/monitor/` or `/var/log/efa/monitor/`. Ensure there are no errors during the generation of the Supportsave. If any issues

arise during the Supportsave generation, try restarting the monitoring service on the active node.

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

Use the **`efa system logging`** commands to view and modify the log level for EFA services. By default, logging is configured with its default parameters and no logging types are excluded. Any log entry exceeding set limits will be lost. This is useful for troubleshooting and validation scenarios that require trace-level logging.

About This Task

By default, the Inventory service log level is set to debugging. You can change the Inventory service logging level from debugging to trace.

Follow this procedure to configure a customized logging.

Note
The trace log level generates additional diagnostic output and should only be used during active troubleshooting or validation. Restore the log level to debugging after validation is complete.

Procedure

1. To configure the log level for a specific service, 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

efa system logging set --service=inventory --level trace
+-----------+------+------+---------------+-------+---------+--------+
|  Service  | Type | Size | Maximum Files | Level | Status  | Reason |
+-----------+------+------+---------------+-------+---------+--------+
| inventory | all  |      |               | trace | Success |        |
+-----------+------+------+---------------+-------+---------+--------+
Logging set
--- Time Elapsed: 122.351032ms ---
```

Note

For information about commands and supported parameters to configure logging, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*

2. Verify the current log level.

   To confirm the log level currently applied to all services, run the **`efa system logging show`** command.

   If you have set the inventory service log level to trace, verify that the newly added trace-level logs are present in the Inventory service logs.
3. Restore the log level to the default.

   Always restore the log level after validation to avoid excessive log output in production environments.

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, 4.1.0*

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.

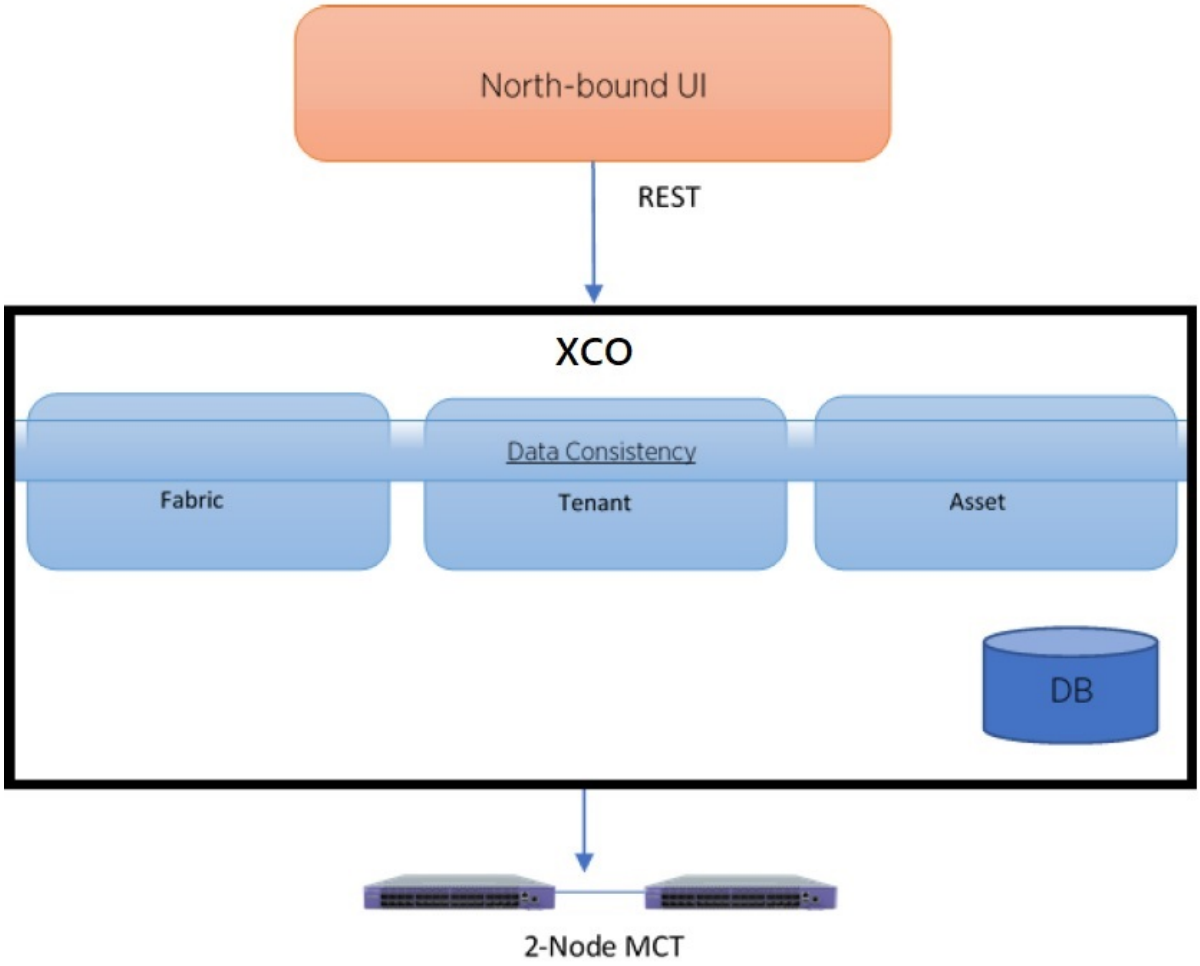

Figure 3: Data consistency overview

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 device configurations. 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.

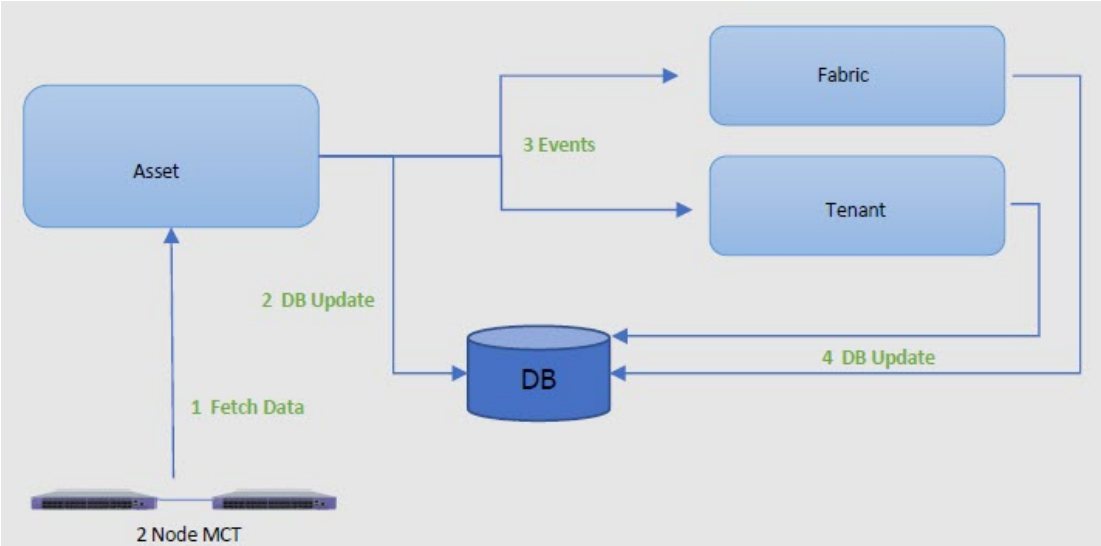

Figure 4: Device discovery and database updates

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

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 reconciled. 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.

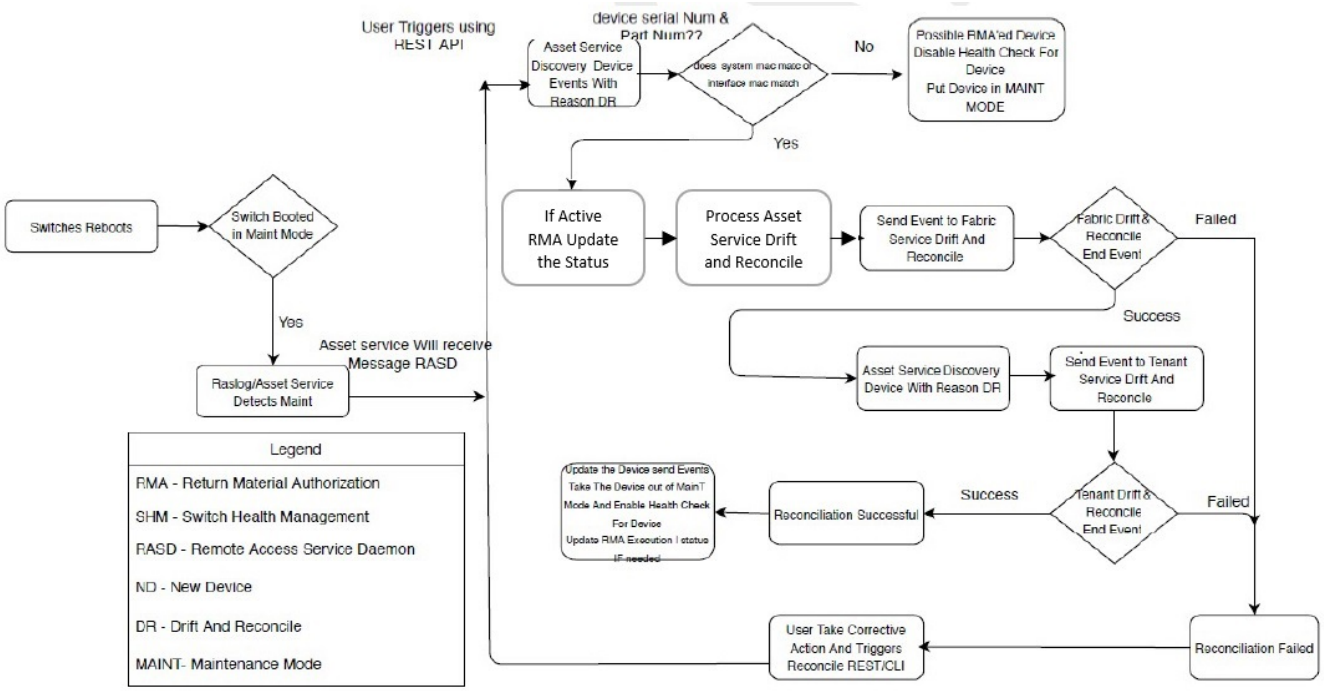

Figure 5: Drift and reconcile workflow

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 EFA 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).

Note

From XCO v4.1.0 onwards, DRC only reconciles interface attributes explicitly configured via XCO.

- Device-default values (discovered during periodic discovery) aren't reconciled
- Prevents unnecessary config pushes post device reboot/RMA.

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 describes the various attributes of an 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, 4.1.0*.

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 setting update --mct-bring-up-delay`
- `efa inventory device setting update --maint-mode-convergence-time`

For more information, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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`

  The **`efa inventory device reload`** command reloads a running SLX device.

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

For more information, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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, 4.1.0*.

SNMP Configuration

The following tables describes 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	
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	

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

Field	Identify Drift	Reconcile Configuration	Idempotency	Notes
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	
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, 4.1.0*.

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 describes 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, 4.1.0*.

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.
Pv4 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 do not 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, 4.1.0*.

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, 4.1.0*.

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, 4.1.0*.

ECMP Max Path

Use the command **`string`** to view route load-sharing maximum paths. Valid values include **`8`**, **`16`**, **`32`**, **`64`** and **`128`** and **`0`** to de-configure.

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, 4.1.0*.

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 a 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, 4.1.0*.

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, 4.1.0*.

Max-Paths Support in SLX Fabric

XCO 4.0.2 introduces BGP maximum-paths support in the SLX fabric, enabling ECMP multipathing for BGP routes across all fabric devices (Leaf, Spine, Super-Spine) based on a fabric-wide setting. All SLX devices in the fabric automatically inherit and enforce the fabric-wide max-paths setting for BGP ECMP with built-in drift detection and consistent multi-rack support.

If the `max_paths = 20`, then all SLX devices will configure maximum-paths 20 under IPv4 Unicast AF.

Key Features

The max-paths feature meets the following functional requirements:

- Support Max-Paths in SLX BGP Underlay

Automatic Config Generation: Applied during initial build & subsequent updates. For every SLX device in the fabric, generate:

```
router bgp
 address-family ipv4 unicast
   maximum-paths <fabric.max_paths>
```

- Use fabric-level setting as the single source of truth

  The **`efa fabric setting update --max-paths <value>`** command updates the database table and triggers config regeneration.
- Store max-paths per device address family

  When setting database table via **BgpIPAf**, ensure you store the applied max-path values correctly:
  - MaxPath: Supported on SLX-OS
  - MaxpathEBgp and MaxpathIBgp: Supported only on Extreme OS ONE

  All set to fabric max-paths
- Drift Detection Required

  Monitors manual deviations from DB-stored values.
- Support Multi-Rack and uniform ECMP

  Ensure consistent ECMP across racks, regardless of link count and topology

The max-paths feature meets the following non-functional requirements:

- No peer-group model changes.
- SLX platforms only - no impact to Extreme OS ONE
- Backward compatible with existing config structures

Data Model Impact

Fabric Setting Table Input

Field	Meaning	Example
Max_paths	Max ECMP paths for BGP underlay	Max_paths = 20

Config Generation Logic

- **Input**
  - Device ID
  - Fabric setting (max_paths)
  - Peer-groups
  - Neighbor table
  - Device role (Leaf or Spine)
- Output (Per Device)
  - Spine Example

    ```
    router bgp
     local-as 64512
     ...
    ```

```
 address-family ipv4 unicast
  maximum-paths 20
  graceful-restart
 !
```

- Leaf Example

```
address-family ipv4 unicast
  network 172.31.254.63/32
  maximum-paths 20
  graceful-restart
!
```

- Render Algorithm (For SLX device only)

```
if device_type == 'SLX':
    maxpaths = fabric_properties.max_paths

    for each AF where AFI == 'ipv4' and SAFI == 'unicast':
        add "maximum-paths <maxpaths>"
        BgpIPAf.MaxPath = maxpaths
        BgpIPAf.MaxpathEBgp = maxpaths
        BgpIPAf.MaxpathIBgp = maxpaths
```

- Handling Defaults

  If `max_paths` is not specified, it defaults to 8 automatically.

Workflow Changes

- **Fabric Creation**
  - `max-paths` defaults to system value (8) unless overridden via CLI.
  - Persisted to DB for reuse.
- **Config Generation**

  Reads max-paths from DB, maps to BgpIPAf.MaxPath fields, and tgeb generates CLI config.
- **Device Provisioning**

  Pushes full BGP config (including maximum-paths) to target devices.
- **Drift Detection**

  Compares the device's running configuration against the expected state stored in the database:
  - If maximum-paths is missing from the running config, drift is detected
  - If the maximum-paths value differs from the DB, drift is detected

Detailed DB to BGP Config Mapping

Key tables driving BGP config generation:

- router_bgps
- router_bgp_address_families
- peer_groups
- remote_neighbor_switch_configs
- fabric_properties

Example: IPv4 Unicast AF Record

```
-- dcapp_fabric.router_bgp_address_families
-----------------------------------------------------
id          = 101
fabric_id   = 3
device_id   = 1
afi         = ipv4
safi        = unicast
is_active   = 1
max_path    = 20
maxpath_ebgp= 20
maxpath_ibgp= 20
```

Generates the following CLI config:

```
router bgp
  address-family ipv4 unicast
    maximum-paths 20
```

Code Changes Overview

1. Read Fabric Setting

   Pull `max_paths` from DB:

   ```
   fabricProps := fabricPropertiesDao.Get(fabricID)
   maxPaths := fabricProps.MaxPaths
   ```

2. Populate BgpIPAf

   ```
   af := BgpIPAf{
       Afi:         "ipv4",
       Safi:        "unicast",
       MaxPath:     int32(maxPaths),
       MaxpathEBgp: int32(maxPaths),
       MaxpathIBgp: int32(maxPaths),
   }
   ```

3. Log Insert Line Under IPv4 AF

   ```
   fmt.Fprintf(&buf, "  maximum-paths %d\n", af.MaxPath)
   ```

4. Initiate Drift Detection

   ```
   if runningConfig.MaxPaths != db.MaxPath {
       addDrift("BGP Max-Paths mismatch")
   }
   ```

Max Path Verification

With `max_paths = 20` in fabric properties:

- **Spine config**: `maximum-paths 20`
- **Leaf config**: `maximum-paths 20`
- **DB table** (`router_bgp_address_families` stores): `max_path = 20`
- Drift detection flags mismatches correctly

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.

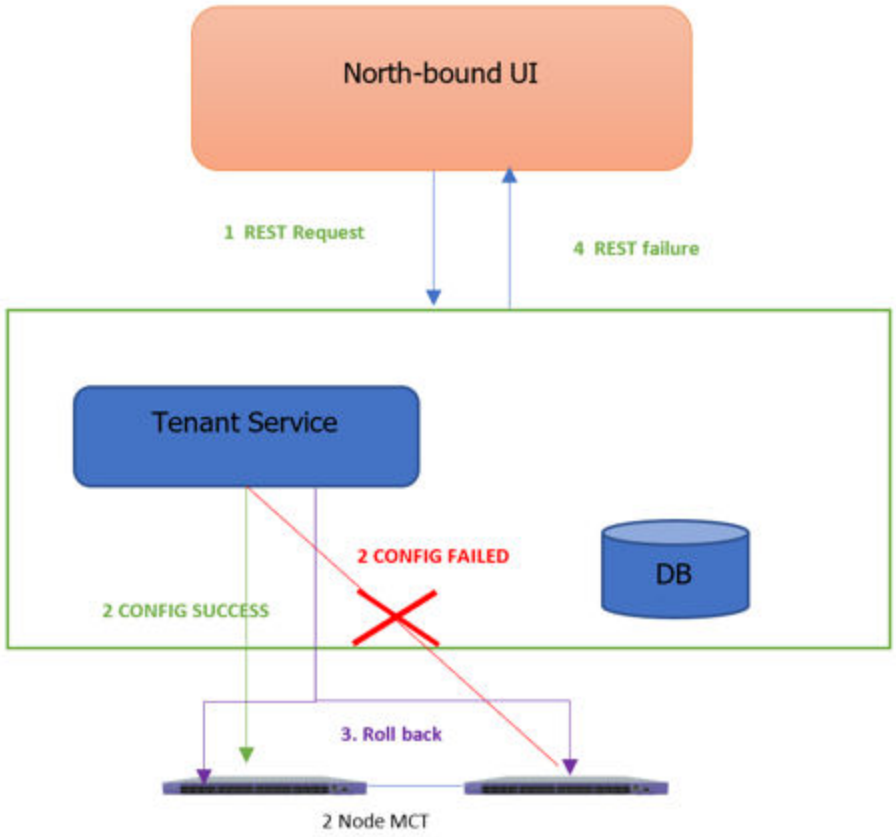

Figure 6: Rollback for failure of one node

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 591.

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.

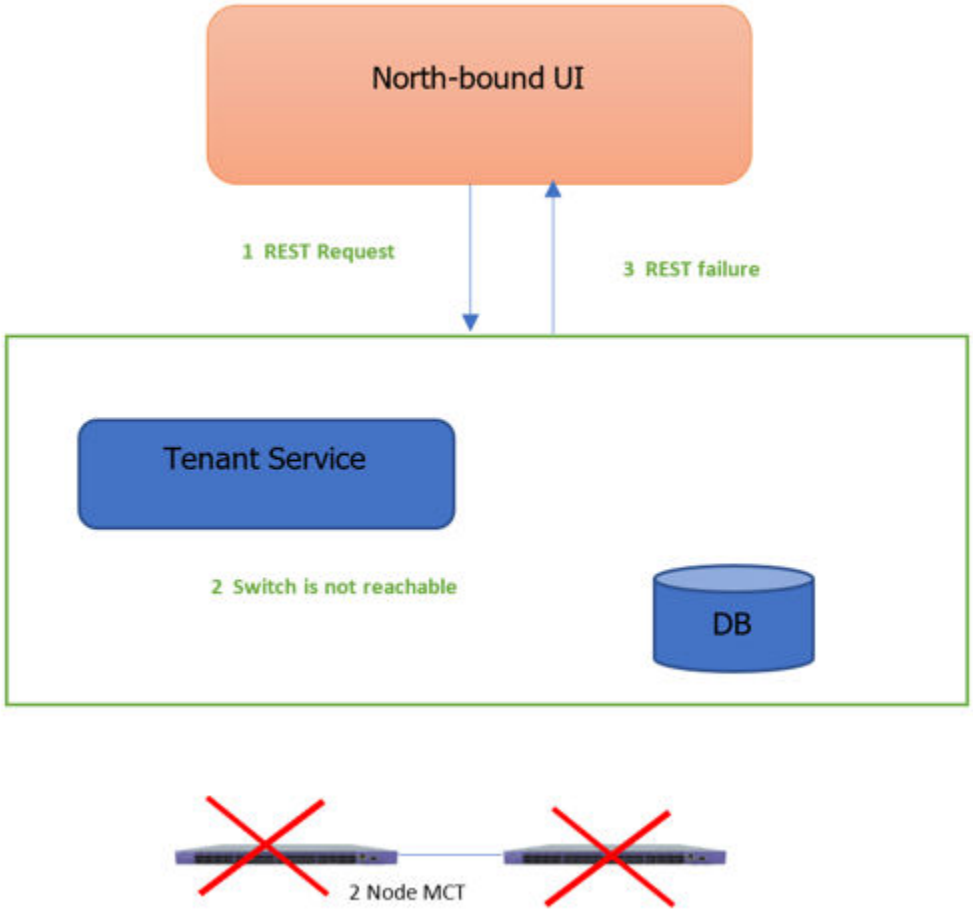

Figure 7: Rollback for failure of both nodes

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, 4.1.0*.

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.

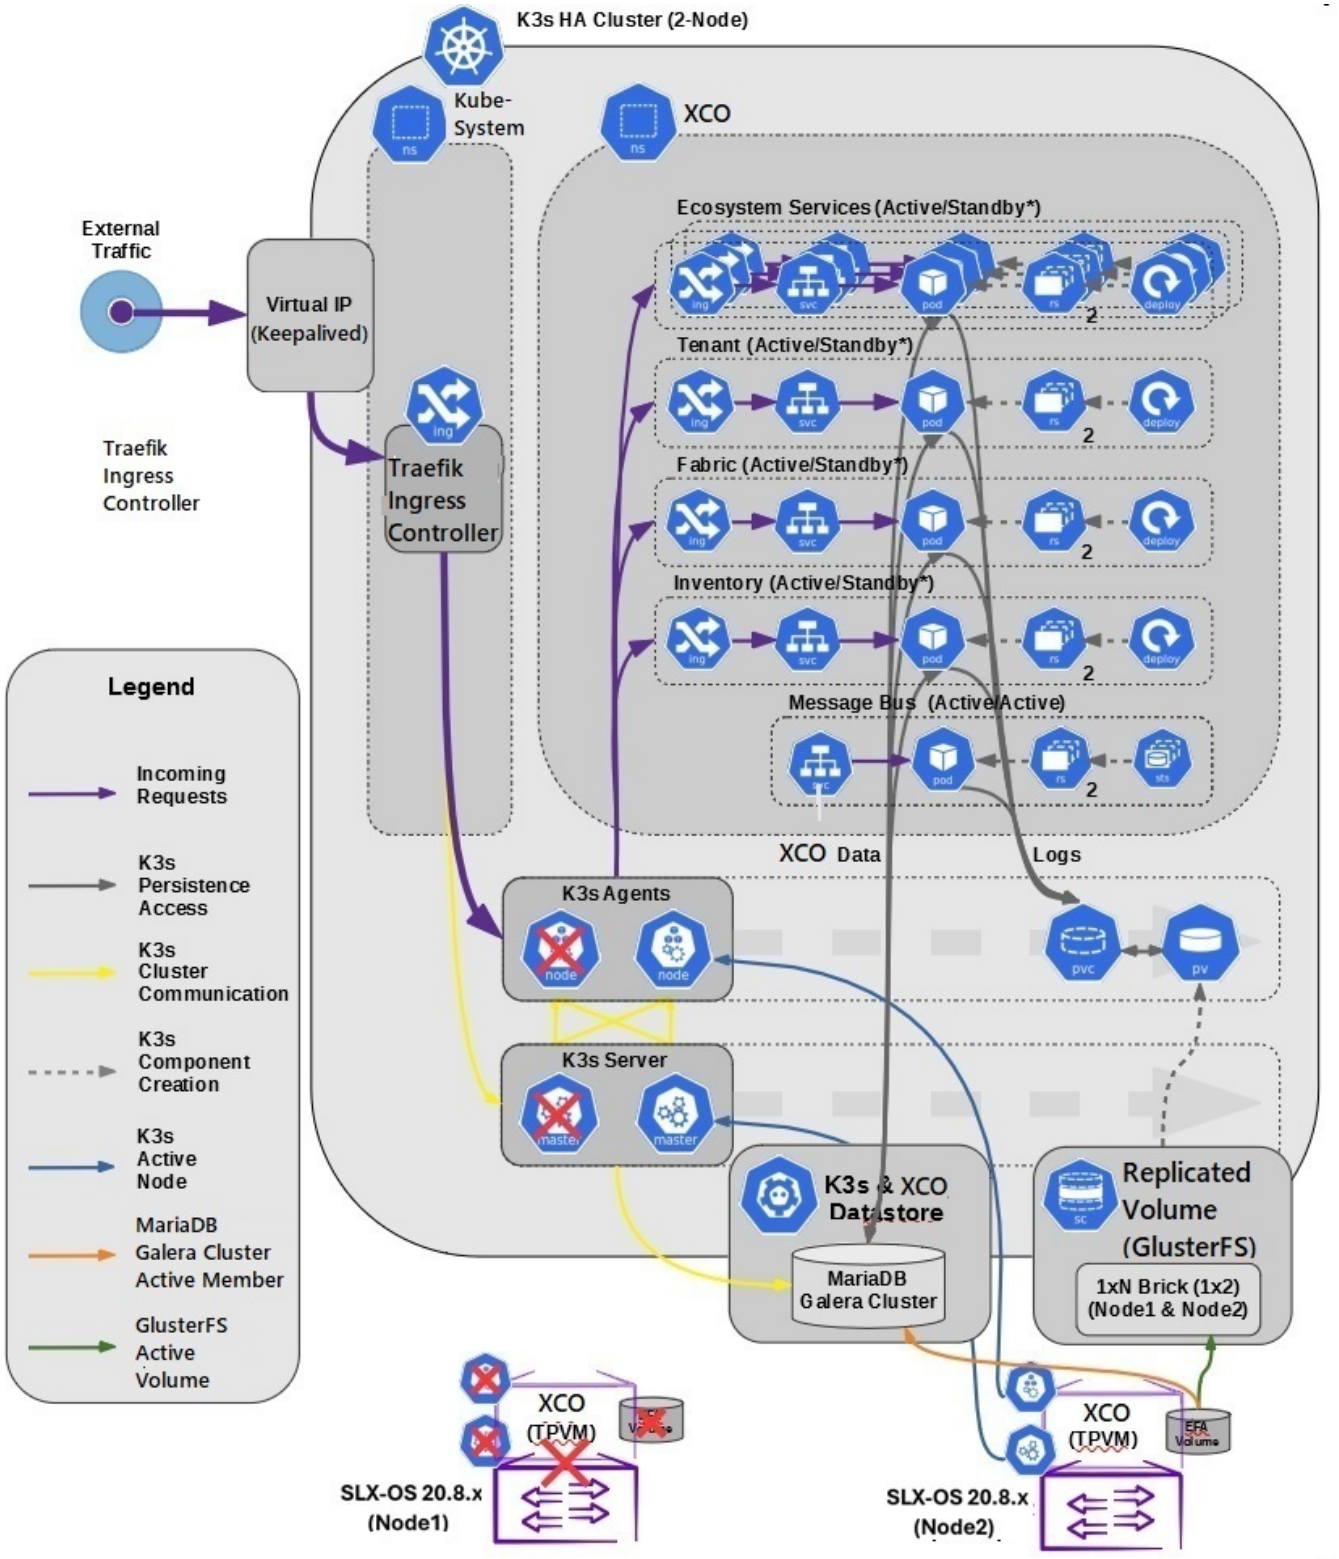

Figure 8: SLX device failure in a two-node cluster

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, 4.1.0*.

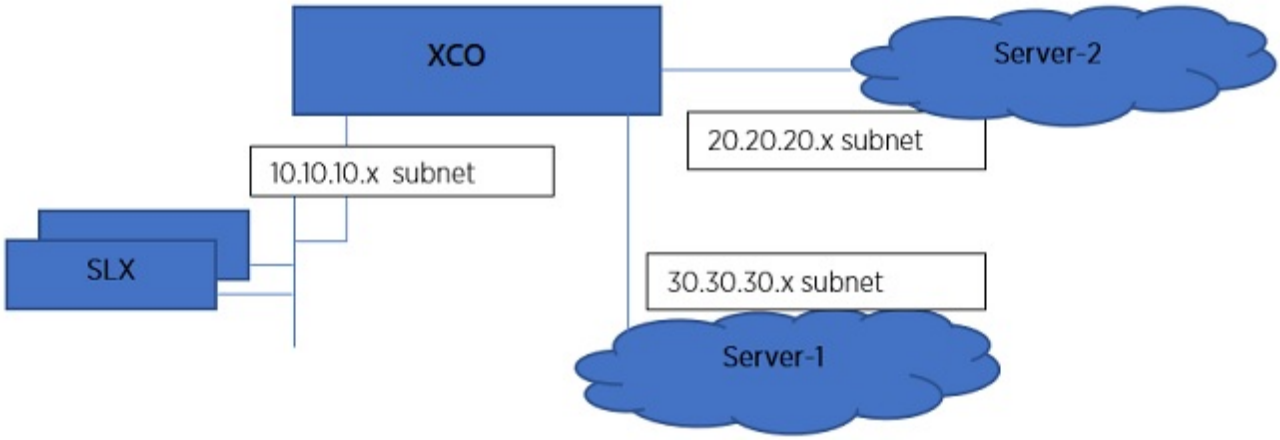

Figure 9: MMIP architecture

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

The create and delete operations do not cause a high-availability (HA) failover. The route is instantiated on the active node of the cluster. When there is a failover from node 1 to node 2, keepalived ensures 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, 4.1.0*.

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
You can add a sub-interface either during installation of XCO or post installation of XCO.

Note
You cannot modify a sub-interface configuration directly, but if changes are needed, you can delete and recreate the sub-interface with the desired values.

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:
   - For IPv4 IP-stack type

     ```
     efa mgmt subinterface staticip add [ --subinterface sub | --ip1 <ipv4-address> |--
     ip2 <ipv4-address> ]
     ```

     Here the syntax shows that ip1 and ip2 will accept only with IPv4 address.

     Note
     The Mix mode configuration with ip1 as IPv4 address and ip2 as IPv6 address or vice versa is not supported.

     ```
     efa mgmt subinterface staticip add -subinterface sub1 --ip1 10.10.10.1/24 -ip2
     10.10.10.2/24
     ```

   - For IPv6 IP-stack type

     ```
     efa mgmt subinterface staticip add [ --subinterface sub | --ip1 <ipv6-address> |--
     ip2 <ipv6-address> ]
     ```

     Here the syntax shows that ip1 and ip2 will accept only with IPv6 address.

     Note
     The Mix mode configuration with ip1 as IPv4 address and ip2 as IPv6 address or vice versa is not supported.

     ```
     efa mgmt subinterface staticip add -subinterface sub1 --ip1 2003::165/64 -ip2
     2003::166/64
     ```

   - For Dual IP-stack type

     ```
     efa mgmt subinterface staticip add [ --subinterface sub | --ip1 <ipv4-address-1> |
     --ip2 <ipv4-address-2> ] --ip3 <ipv6-address-1> | --ip4 <ipv6-address-2>]
     ```

     Here the syntax shows that ip1 and ip2 addresses will be available only to IPv4 addresses, and that ip3 and ip4 addresses will be available only to IPv6 addresses.

     To set up an optional static IPv6 address, ensure that the sub-interface to which it is assigned already has an IPv6 address configured. All other assumptions and validations remain consistent with those used for static IPv4 addresses in the IPv4 stack.

     ```
     efa mgmt subinterface staticip add -subinterface sub1 --ip1 10.10.10.1/24 --ip2
     10.10.10.2/24 -ip3 2000::1/64 --ip4 2000::2/64
     ```

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 for the IPv4 IP-Stack.

You can use this procedure for Dual IP-Stack type with Dual static IP configuration. For more information, see Configure Static IP Addresses for Management Sub Interfaces on page 148.

This procedure does not affect 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
   ```

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.

Use this procedure when you do not configure DNS server before XCO installation. When you update DNS servers on the host, use the script to update the same on XCO services.

Note

- This procedure should be used when the user does not configure dns server before XCO installation. When the user updates dns servers on the host, this script has to be used to update the same on XCO services.
- 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
```

Note

You can change the DNS IP through netplan using the `update-dns.sh` script. For more details, see *ExtremeCloud Orchestrator Release Notes, 4.1.0*.

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.
```

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.

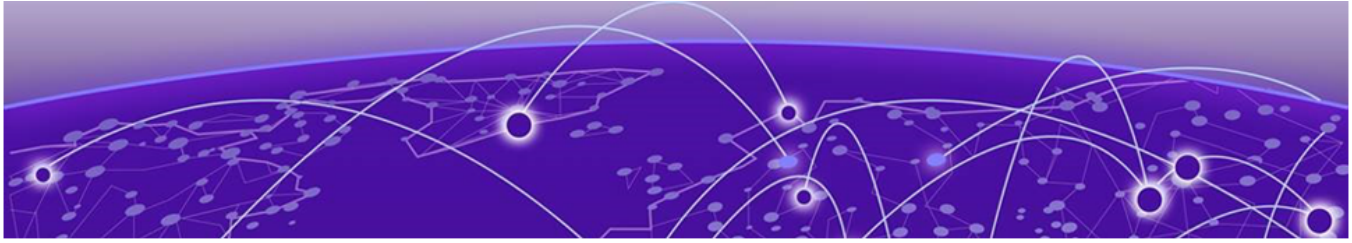

Fabric Infrastructure Provisioning

Learn about automating the fabric underlay and overlay configuration.

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; 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, 4.1.0*.
- 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, 4.1.0*.
- Fabric links (links between spine and leaf, spine and super-spine, and super-spine and border leaf) cannot be a breakout port.
- 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 the configuration steps of each platform.

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.

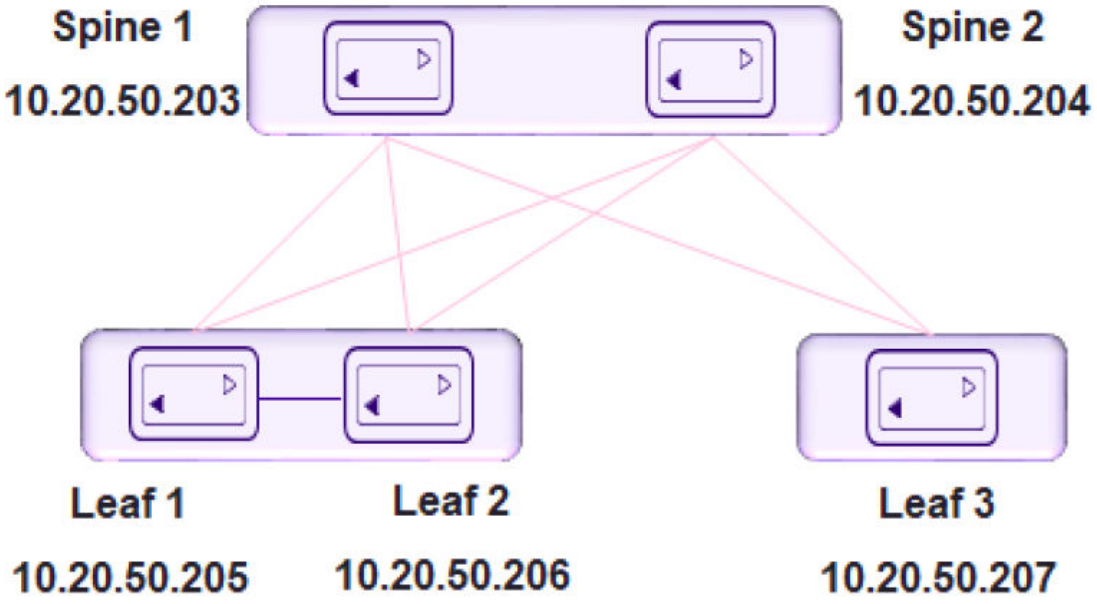

5-Stage Clos

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.

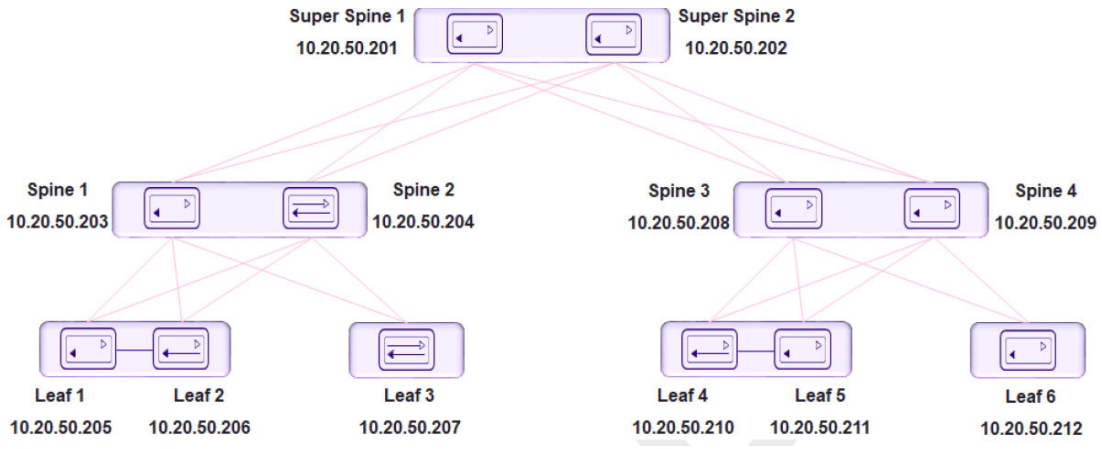

Configure a 3-Stage Clos Fabric

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 a fabric.

```
efa fabric create
```

The following example creates a fabric:

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

2. Add a device to the fabric.

```
efa fabric device add
```

The following 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 <username> --password
<password>
```

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.

   The **`efa inventory device interface list -ip <device-ip>`** displays the list of interfaces and details for the specified IP address, including the application state that indicates whether the device configuration is synchronized with XCO or has drifted (refreshed or deleted).

   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
```

The following example configures the fabric:

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

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`.

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.

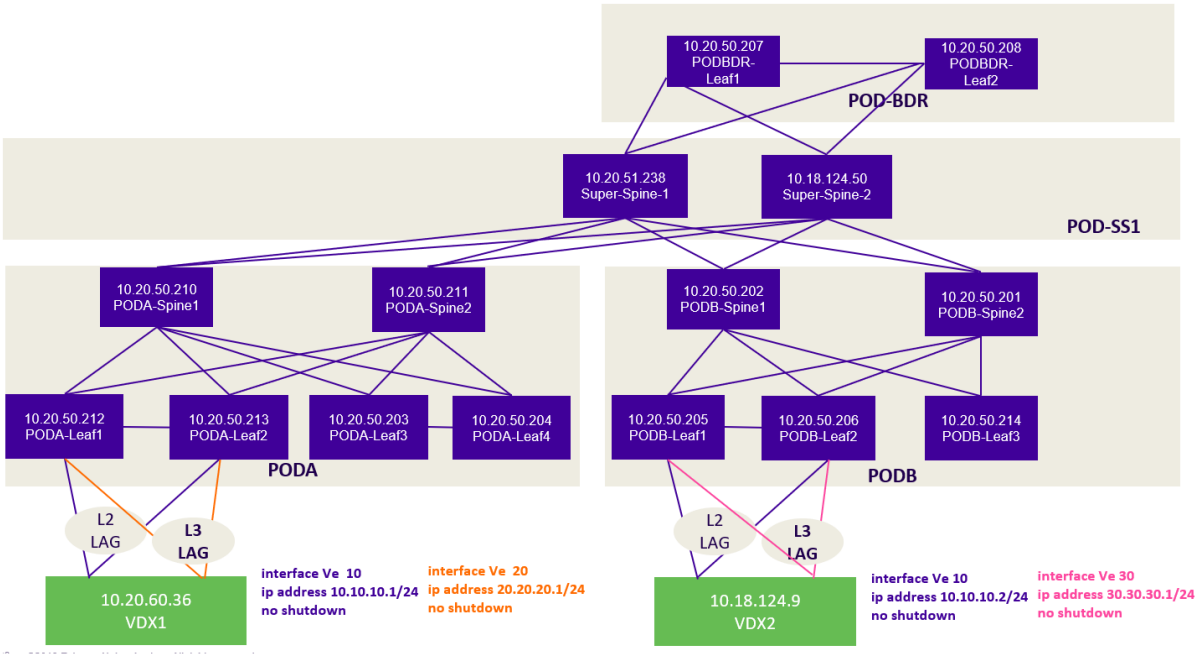

Figure 10: 5-Stage Clos fabric topology

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
   ```

   The following example creates the fabric:

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

2. Add a device to the fabric.

   ```
   efa fabric device add
   ```

   The following example adds a device to the fabric:

   ```
   efa fabric device add--name stage5 --username <username> --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
   ```

   The following example adds multiple devices to the fabric:

   ```
   fa fabric device add-bulk --name stage5 --username <username> --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
   ```

   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
```

The following example configures the fabric topology:

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

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.

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.

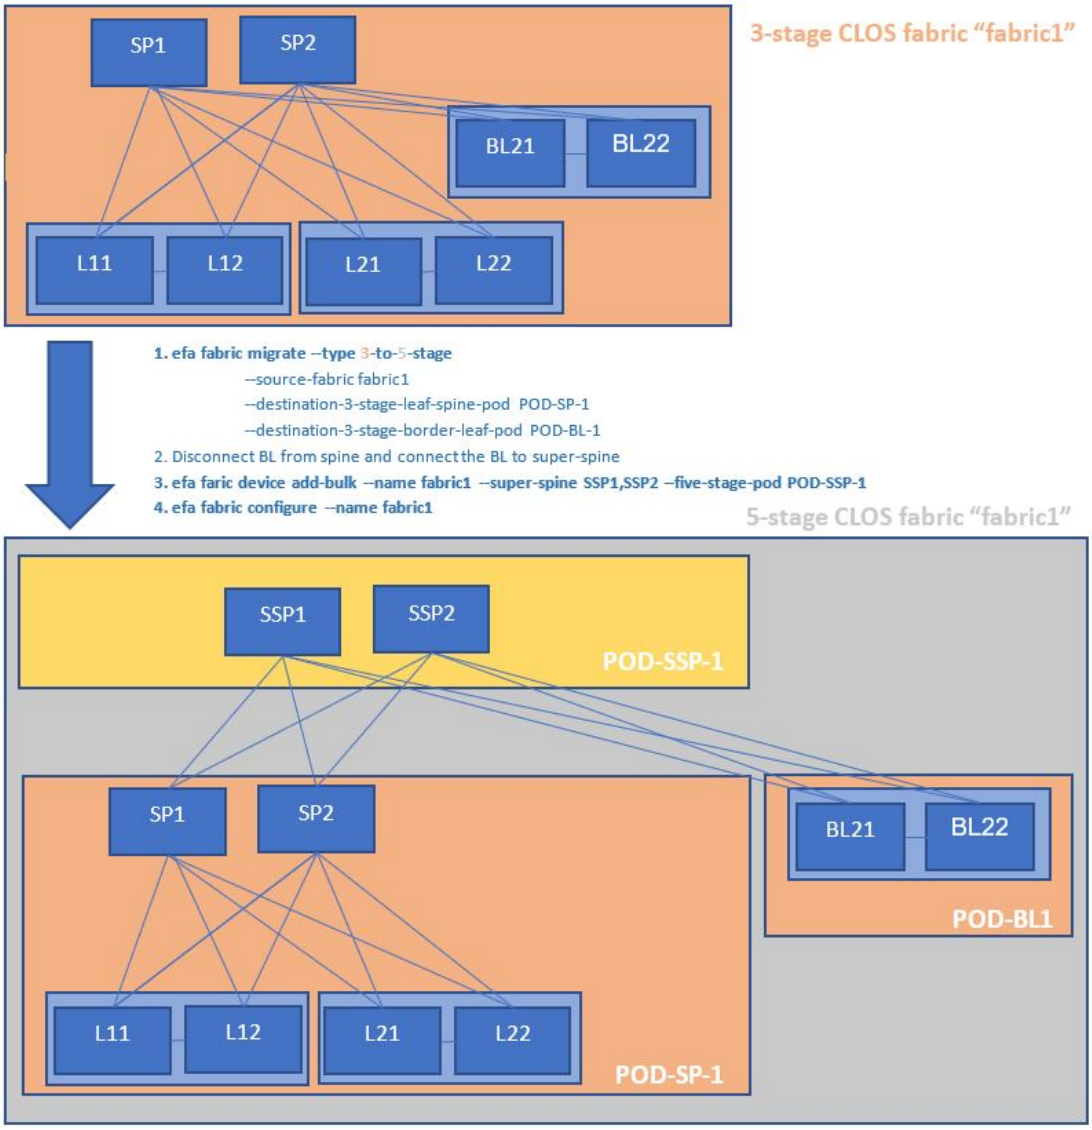

Migrate a 3-Stage Clos to 5-Stage Clos Fabric

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 163
2. Migrate a 3-Stage Clos to 5-Stage Clos Fabric on page 163
3. Disconnect Border-leafs from Spine and Connect to Super-spine on page 164
4. Addition of Super-spine Devices to the Migrated 5-stage Clos Fabric on page 165
5. Configure Migrated 5-stage Clos Fabric on page 165
6. Traffic Disruption during Fabric Configure on page 165

7. Verification of Fabric Underlay Configuration on the Migrated 5-stage Clos Fabric on page 166
8. Verification of Fabric Physical Underlay and Overlay Topology on the Migrated 5-stage Clos Fabric on page 169

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 <username> --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:

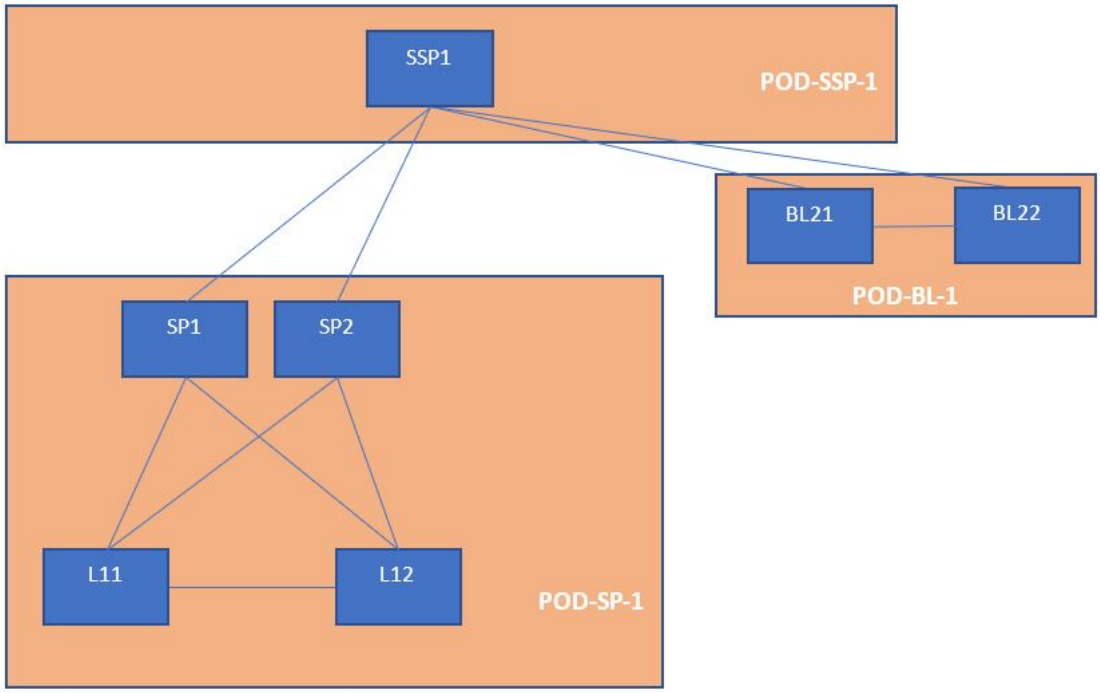

Disconnect Border-leafs from Spine and Connect to Super-spine

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 <username> --password <password>
--super-spine 10.17.112.228   --five-stage-pod POD-SSP-1
```

Configure 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

```
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
```

Super Spine

```
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-
```

Boarder 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 Extreme 8520 and SLX 9140 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 Extreme 8520, Extreme 8720, 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

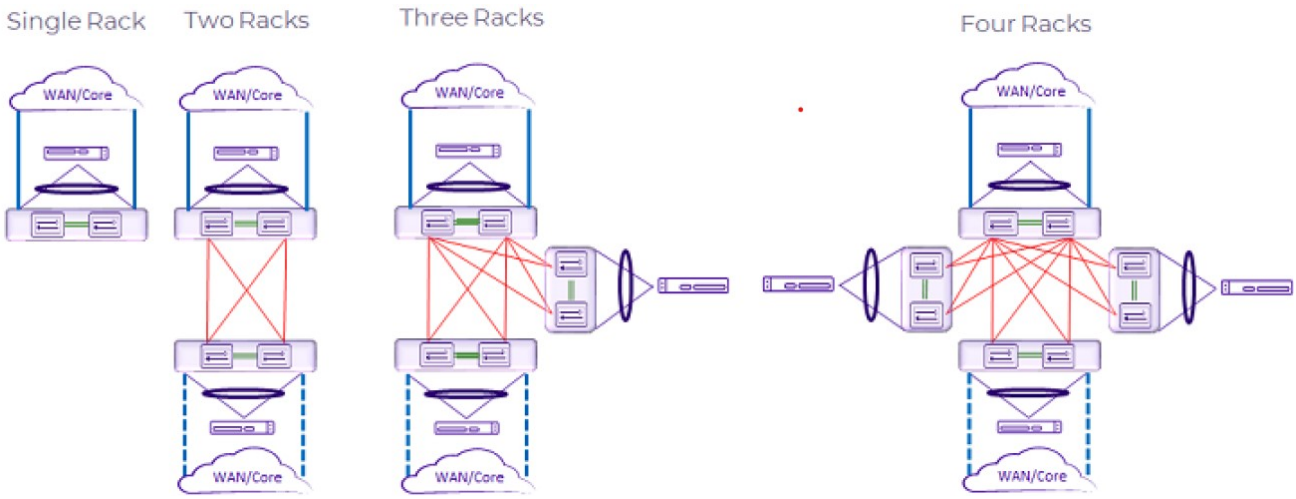

Figure 11: Supported small data center topologies

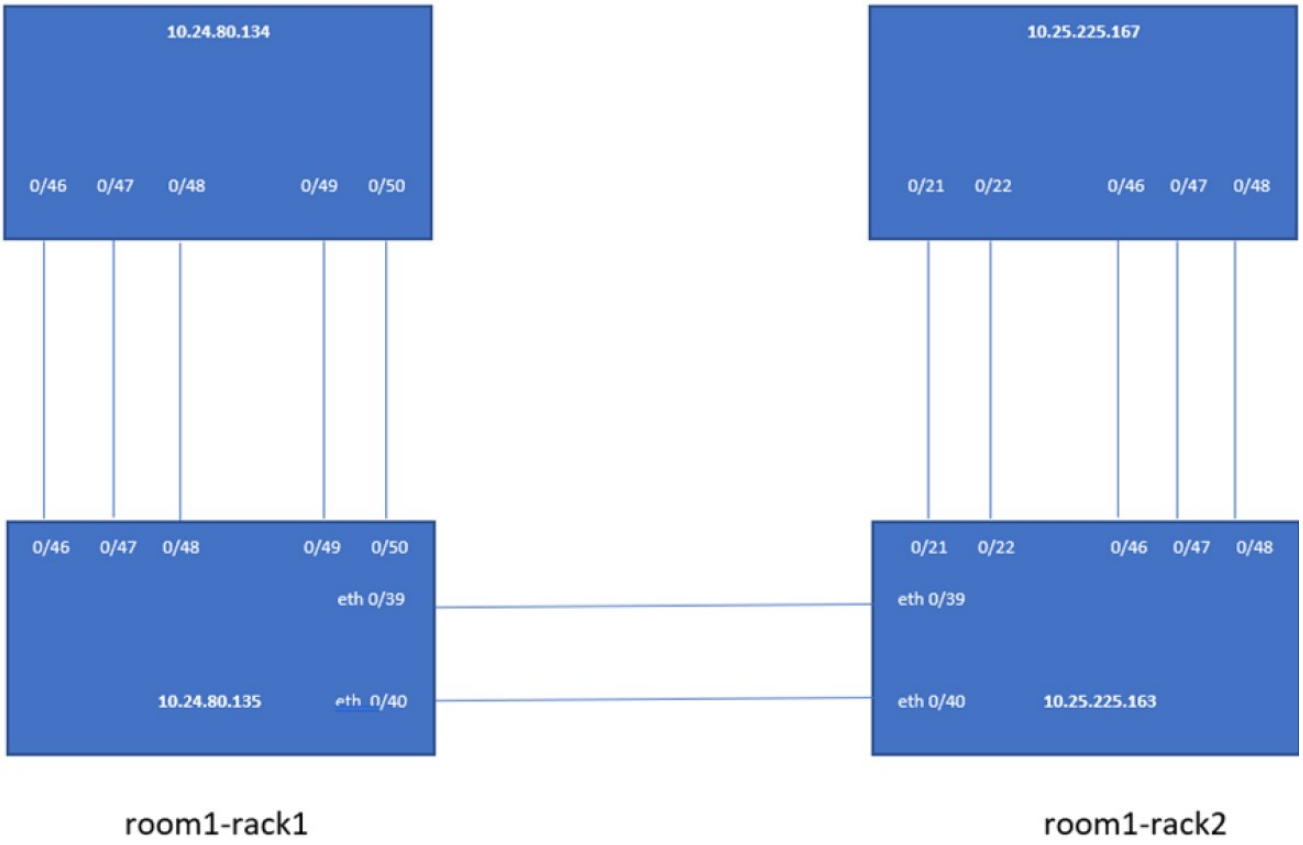

Figure 12: Multi-rack configuration example

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
```

The following example creates a fabric:

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

2. Add a device to the fabric.

```
efa fabric device add
```

The following example adds a device to the fabric:

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

The following 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
```

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
```

This example configures the fabric.

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

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.

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 centers 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, that 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.

Static ICL in Small Data Center

XCO supports both Static and Dynamic ICL.

Prior to EFA 2.6.0, ports were defined statically to be part of the ICL or MCT port-channel (**Static ICL**) using the following fabric settings for a non-Clos fabric:

- rack-mct-ports (Applicable for the racks containing high density ports)
- rack-ld-mct-ports (Applicable for the racks containing low density ports)

The fabric settings had the default values which could be modified based on the inter-connectivity between the MCT devices.

From EFA 2.6.0 onwards, the fabric settings, which were used to define the static ICL ports, are removed from XCO to support Dynamic ICL. In a dynamic ICL, all the ports inter-connecting the MCT devices will automatically participate in the ICL or MCT port-channel.

The Dynamic ICL is the **default** behavior. The static ICL support is applicable only for a non-Clos fabric.

For the fabrics created prior to XCO 3.3.0, the Dynamic ICL will continue to be effective. You can transition to static ICL by modifying the `rack-mct-scheme`, `rack-mct-ports`, and `rack-ld-mct-ports` fabric settings.

For the fabrics created from XCO 3.3.0 onwards, the **default** behavior is Dynamic ICL unless you choose to provide the `rack-mct-scheme`, `rack-mct-ports`, and `rack-ld-mct-ports` fabric settings.

When you choose static ICL scheme and the ports inter-connecting the MCT devices do not contain all the ports provided in the `rack-mct-ports` and `rack-ld-mct-ports` fabric

settings, system shows errors in the “efa fabric show”, “efa fabric error show”, and the “efa fabric health show” output when you add devices in fabric (fabric device add).

Note

- The fabric settings `rack-mct-ports` and `rack-ld-mct-ports` are applicable to all the high density and low density racks respectively.
- The fabric setting `rack-mct-ports` is applicable for the platforms with high density ports (devices with higher number of ports) and `rack-ld-mct-ports` is applicable for the platforms with low density platforms (devices with lower number of ports).
- Hardware platforms categorized as low and high density are described in the following table:

Low density hardware platforms	High density hardware platforms
SLX 9640, SLX 9740_40C, SLX 9740_80C, Extreme 8720, Extreme 8820_40C, Extreme 8820_80C	SLX 9540, Extreme 8730, Extreme 8740, Extreme 8520, and Extreme 8520T

Configure Static and Dynamic ICL

You can configure static and dynamic ICL.

About This Task

Follow this procedure to configure static or dynamic ICL.

Procedure

To configure static or dynamic ICL, run the following command:

```
efa fabric setting update --name <fabric-name>
          --rack-mct-scheme <static | dynamic>
          --rack-mct-ports <port-list>
          --rack-ld-mct-ports <port-list>
```

Example

- The following is an example output of dynamic ICL configuration:

```
Rack1-Device1(config)# do show lldp neighbors | include Rack1-Device2
Eth 0/13    120    119    Ethernet 0/13     Eth 0/13    f46e.95a0.c805   77822     24     Rack1-
Device2
Eth 0/14    120    118    Ethernet 0/14     Eth 0/14    f46e.95a0.c805   77825     23     Rack1-
Device2
Eth 0/15    120    118    Ethernet 0/15     Eth 0/15    f46e.95a0.c805   77822     152    Rack1-
Device2
Rack1-Device1(config)#

Rack1-Device2(config)# do show lldp neighbors | include Rack1-Device1
Eth 0/13    120    110    Ethernet 0/13     Eth 0/13    f46e.95a2.b805   92706     23     Rack1-
Device1
Eth 0/14    120    110    Ethernet 0/14     Eth 0/14    f46e.95a2.b805   92699     23     Rack1-
Device1
Eth 0/15    120    110    Ethernet 0/15     Eth 0/15    f46e.95a2.b805   92697     152    Rack1-
Device1
Rack1-Device2(config)#

efa fabric create --name fabric1 --type non-clos
```

```
(efa:root)root@administrator:~# efa fabric setting show --name fabric1 --advanced | grep -i  mct
| MCT Link IP Range              | 10.20.20.0/24          |
| MCT PortChannel                | 64                     |
| Rack MCT Scheme                | Dynamic                |
| Rack MCT Ports                 |                        |

efa fabric device add-bulk --name fabric1 --ip 10.20.246.1-2 --rack rack1
Add Device(s) [Success]

        Addition of Leaf device with ip-address = 10.20.246.2 [Succeeded]
        Addition of Leaf device with ip-address = 10.20.246.1 [Succeeded]
Validate Fabric [Success]

efa fabric configure --name fabric1
```

```
Rack1-Device1(config)# do show lldp neighbors | include Rack1-Device2
Eth 0/13    120    119    Ethernet 0/13     Eth 0/13     f46e.95a0.c805   77822    24   Rack1-
Device2
Eth 0/14    120    118    Ethernet 0/14     Eth 0/14     f46e.95a0.c805   77825    23   Rack1-
Device2
Eth 0/15    120    118    Ethernet 0/15     Eth 0/15     f46e.95a0.c805   77822    152  Rack1-
Device2
Rack1-Device1(config)#

Rack1-Device1(config)# do show port-channel 64
 LACP Aggregator: Po 64
 Aggregator type: Standard
  Admin Key: 0064 - Oper Key 0064
  Partner System ID - 0x8000,f4-6e-95-a0-c8-05
  Partner Oper Key 0064
 Flag * indicates: Primary link in port-channel
 Number of Ports: 3
 Minimum links: 1
 Member ports:
   Link: Eth 0/13 (0xC01A100) sync: 1
   Link: Eth 0/14 (0xC01C100) sync: 1   *
   Link: Eth 0/15 (0xC01E100) sync: 1
Rack1-Device1(config)#
```

```
Rack1-Device2(config)# do show lldp neighbors | include Rack1-Device1
Eth 0/13     120    110    Ethernet 0/13     Eth 0/13    f46e.95a2.b805   92706     23    Rack1-
Device1
Eth 0/14     120    110    Ethernet 0/14     Eth 0/14    f46e.95a2.b805   92699     23    Rack1-
Device1
Eth 0/15     120    110    Ethernet 0/15     Eth 0/15    f46e.95a2.b805   92697     152   Rack1-
Device1
Rack1-Device2(config)#

Rack1-Device2(config)# do show port-channel 64
 LACP Aggregator: Po 64
 Aggregator type: Standard
  Admin Key: 0064 - Oper Key 0064
  Partner System ID - 0x8000,f4-6e-95-a2-b8-05
  Partner Oper Key 0064
 Flag * indicates: Primary link in port-channel
 Number of Ports: 3
 Minimum links: 1
 Member ports:
   Link: Eth 0/13 (0xC01A100) sync: 1   *
   Link: Eth 0/14 (0xC01C100) sync: 1
   Link: Eth 0/15 (0xC01E100) sync: 1
Rack1-Device2(config)#
```

- The following is an example output of static ICL configuration:

```
Rack1-Device1(config)# do show lldp neighbors | include Rack1-Device2
Eth 0/13    120     119   Ethernet 0/13    Eth 0/13    f46e.95a0.c805   77822    24   Rack1-Device2
```

```
Eth 0/14    120     118   Ethernet 0/14    Eth 0/14    f46e.95a0.c805   77825    23   Rack1-Device2
Eth 0/15    120     118   Ethernet 0/15    Eth 0/15    f46e.95a0.c805   77822    152  Rack1-Device2
Rack1-Device1(config)#

Rack1-Device2(config)# do show lldp neighbors | include Rack1-Device1
Eth 0/13    120    110    Ethernet 0/13    Eth 0/13    f46e.95a2.b805   92706    23   Rack1-Device1
Eth 0/14    120    110    Ethernet 0/14    Eth 0/14    f46e.95a2.b805   92699    23   Rack1-Device1
Eth 0/15    120    110    Ethernet 0/15    Eth 0/15    f46e.95a2.b805   92697    152  Rack1-Device1
Rack1-Device2(config)#

(efa:root)root@administrator:~# efa fabric create --name fabric1 --type non-clos

(efa:root)root@administrator:~# efa fabric setting show --name fabric1 --advanced | grep -i  mct

(efa:root)root@administrator:~# efa fabric setting show --name fabric1 --advanced | grep -i  mct
| MCT Link IP Range              | 10.20.20.0/24          |
| MCT PortChannel                | 64                     |
| Rack MCT Scheme                | Dynamic                |
| Rack MCT Ports                 |                        |
(efa:root)root@administrator:~#

(efa:root)root@administrator:~# efa fabric setting update --name fabric1 --rack-mct-scheme static
--rack-mct-ports 0/31,0/32

(efa:root)root@administrator:~# efa fabric setting show --name fabric1 --advanced | grep -i  mct
| MCT Link IP Range              | 10.20.20.0/24          |
| MCT PortChannel                | 64                     |
| Rack MCT Scheme                | Static                 |
| Rack MCT Ports                 | 0/31,0/32              |
(efa:root)root@administrator:~#

(efa:root)root@administrator:~# efa fabric device add-bulk --name fabric1 --ip 10.20.246.1-2 --
rack rack1 --username <username> --password <password>
Add Device(s) [Success]

        Addition of Leaf device with ip-address = 10.20.246.1 [Succeeded]
        Addition of Leaf device with ip-address = 10.20.246.2 [Succeeded]
Validate Fabric [Failed]
        Config MisMatch
        MCT Cluster Config missing between rack devices 10.20.246.1 and 10.20.246.2.Please remove
rack devices and re-add.
        Missing Links
        Device 10.20.246.1 is not connected to device 10.20.246.2 on Mct ports 0/31,0/32
Error :  fabric validation failed

(efa:root)root@administrator:~# efa fabric device remove --name fabric1 --ip 10.20.246.1-2

(efa:root)root@administrator:~# efa fabric setting update --name fabric1 --rack-mct-scheme static
--rack-mct-ports 0/13,0/14

(efa:root)root@administrator:~# efa fabric setting show --name fabric1 --advanced | grep -i  mct
| MCT Link IP Range              | 10.20.20.0/24          |
| MCT PortChannel                | 64                     |
| Rack MCT Scheme                | Static                 |
| Rack MCT Ports                 | 0/13,0/14              |
(efa:root)root@administrator:~#

(efa:root)root@administrator:~# efa fabric device add-bulk --name fabric1 --ip 10.20.246.1-2 --
rack rack1

(efa:root)root@administrator:~# efa fabric configure --name fabric1

Rack1-Device1(config)# do show lldp neighbors | include Rack1-Device2
Eth 0/13   120    119    Ethernet 0/13    Eth 0/13    f46e.95a0.c805   77822    24    Rack1-Device2
Eth 0/14   120    118    Ethernet 0/14    Eth 0/14    f46e.95a0.c805   77825    23    Rack1-Device2
Eth 0/15   120    118    Ethernet 0/15    Eth 0/15    f46e.95a0.c805   77822    152   Rack1-Device2
```

```
Rack1-Device1(config)#

Rack1-Device1(config)# do show port-channel 64
 LACP Aggregator: Po 64
 Aggregator type: Standard
  Admin Key: 0064 - Oper Key 0064
  Partner System ID - 0x8000,f4-6e-95-a0-c8-05
  Partner Oper Key 0064
 Flag * indicates: Primary link in port-channel
 Number of Ports: 2
 Minimum links: 1
 Member ports:
   Link: Eth 0/13 (0xC01A100) sync: 1
   Link: Eth 0/14 (0xC01C100) sync: 1   *

Rack1-Device1(config)#

Rack1-Device2(config)# do show lldp neighbors | include Rack1-Device1
Eth 0/13      120    110   Ethernet 0/13    Eth 0/13     f46e.95a2.b805   92706     23   Rack1-
Device1
Eth 0/14      120    110   Ethernet 0/14    Eth 0/14     f46e.95a2.b805   92699     23   Rack1-
Device1
Eth 0/15      120    110   Ethernet 0/15    Eth 0/15     f46e.95a2.b805   92697     152  Rack1-
Device1
Rack1-Device2(config)#

Rack1-Device2(config)# do show port-channel 64
 LACP Aggregator: Po 64
 Aggregator type: Standard
  Admin Key: 0064 - Oper Key 0064
  Partner System ID - 0x8000,f4-6e-95-a2-b8-05
  Partner Oper Key 0064
 Flag * indicates: Primary link in port-channel
 Number of Ports: 2
 Minimum links: 1
 Member ports:
   Link: Eth 0/13 (0xC01A100) sync: 1   *
   Link: Eth 0/14 (0xC01C100) sync: 1

Rack1-Device2(config)#
```

Static ICL to Dynamic ICL Conversion

You can convert a static ICL into a dynamic ICL.

You can update the fabric settings `rack-mct-scheme` on an already deployed fabric provided all the racks of the existing fabric have same ports participating in the ICL or MCT port-channel.

Ensure that the ICL or MCT LLDP ports of all the racks match exactly with the `rack-mct-ports` and `rack-Id-mct-ports` fabric setting.

Shutdown or disable (lldp disable) the additional LLDP ports followed by `fabric configure`, and then re-attempt the static ICL to dynamic ICL conversion. Once the conversion is completed, bring up the additional LLDP links followed by `fabric configure`.

Note
When you modify the fabric setting `rack-mct-scheme` to "dynamic", the fabric setting `rack-Id-mct-ports` is reset to the empty string.

Dynamic ICL to Static ICL Conversion

You can convert a dynamic ICL into a static ICL.

You can update the fabric settings `rack-mct-scheme`, `rack-mct-ports`, and `rack-ld-mct-ports` (with the definitive set of ports) on an already deployed fabric provided all the racks of the existing fabric have same ports participating in the ICL or MCT port-channel.

Ensure that the ICL or MCT LLDP ports of all the racks match exactly with the `rack-mct-ports` and `rack-ld-mct-ports` fabric setting.

Shutdown or disable (or lldp disable) the additional LLDP ports followed by `fabric configure`, and then re-attempt the dynamic ICL to static ICL conversion. Once the conversion is completed, bring up the additional LLDP links.

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, 4.1.0*.

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 <username> --password <password> --name CNCF`
Add the second device	`efa fabric device add --ip 10.24.80.159 --hostname slx-b --rack pod1 --username <username> --password <password> --name CNCF`
Configure the fabric	`efa fabric configure --name CNCF`

View Device Error in Clos and Non-Clos Fabric

Use the REST APIs or CLIs to view the errors in Clos and non-Clos fabric devices.

About This Task

Follow this procedure to show errors in Clos and non-Clos fabric devices.

Procedure

To show the device error or failure reason, do any of the following:

- Run the **`efa fabric error`** show command.

- Use the `GET /fabric/fabrics/errors` REST API.

Note
Use the `--name [fabric_name]` option to filter down the command output to the fabric name.

Example

```
$ efa fabric error show
+-------------+--------------+-------+------------+---------------------------------+
| FABRIC NAME |  IP ADDRESS  | ROLE  | ERROR TYPE |          ERROR REASON           |
+-------------+--------------+-------+------------+---------------------------------+
| fab-non-clos| NA           | NA    | clos error | No Devices                      |
| fab-clos    | 10.20.246.12 | Leaf  | clos error | Leaf Device 10.20.246.12 not    |
|             |              |       |            | connected to Spine Device       |
|             |              |       |            | 10.20.246.19                    |
| fab-clos    | 10.20.246.12 | Leaf  | clos error | Leaf Device 10.20.246.12 not    |
|             |              |       |            | connected to Spine Device       |
|             |              |       |            | 10.20.246.20                    |
| fab-clos    | 10.20.246.19 | Spine | clos error | Spine Device 10.20.246.19       |
|             |              |       |            | not connected to Leaf Device    |
|             |              |       |            | 10.20.246.12                    |
| fab-clos    | 10.20.246.20 | Spine | clos error | Spine Device 10.20.246.20       |
|             |              |       |            | not connected to Leaf Device    |
|             |              |       |            | 10.20.246.12                    |
| fab-clos    | 10.20.246.19 | Spine | clos error | Spine Device 10.20.246.19       |
|             |              |       |            | connected to Spine Device       |
|             |              |       |            | 10.20.246.20                    |
| fab-clos    | 10.20.246.20 | Spine | clos error | Spine Device 10.20.246.20       |
|             |              |       |            | connected to Spine Device       |
|             |              |       |            | 10.20.246.19                    |
+-------------+--------------+-------+------------+---------------------------------+
```

Router ID and VTEP Loopback IP Allocation in Clos and Non-Clos Fabric

You can choose uniform or granular IP allocation scheme for Router ID and VTEP Loopback.

From XCO 3.3.0 onwards, the default scheme for an IP allocation is **Uniform**. This means that the IP allocation scheme for both Router ID Loopback and VTEP Loopback remains unchanged. An IP is allocated from the IP range defined in the `loopback-ip-range` fabric setting.

To use different IP range for Router ID Loopback and VTEP Loopback, choose the **Granular** IP allocation scheme. Provide different IP range for Router ID Loopback and VTEP Loopback using the `router-id-loopback-ip-range` and `vtep-loopback-ip-range` fabric settings.

The IP range provided by you should not overlap with the IP ranges provided in other fabric settings. For example, P2P Link IP range and MCT Link IP range.

Note

- The existing fabrics prior to XCO 3.3.0 and the newly created fabrics from XCO 3.3.0 onwards continue to operate in the older (uniform) scheme by default. You must explicitly choose the newer (granular) scheme .
- You cannot dynamically change the older scheme to newer scheme if the devices are already added to the fabric.
- Once the fabric is configured, you can only expand the IP Range. Compressing the IP range on an active fabric is not supported.

Allocate IP using Uniform Loopback Scheme

You can choose uniform IP allocation scheme for Router ID and VTEP Loopback.

About This Task

Follow this procedure to allocate an IP for Router ID and VTEP Loopback using uniform scheme.

Procedure

To allocate an IP for Router ID and VTEP Loopback using uniform scheme, run the following command:

```
efa fabric setting update --name <fabric-name> --loopback-scheme <uniform|granular>
efa fabric setting update --name <fabric-name> --loopback-ip-range <ip-range>
```

Example

```
Fabric Name: fs-1, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-success,
Fabric Health: Green
+-------------+------+-----------+------------+------+-------------+-------------+------------
+---------+------+----+
| IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE | DEVICE STATE
|  APP STATE  | CONFIG     | PENDING | VTLB | LB |
|             |      |           |            |      |             |             |
GEN REASON | CONFIGS | ID   | ID |
+-------------+------+-----------+------------+------+-------------+-------------+------------
+---------+------+----+
| 10.20.246.3 | r1   | NH-Leaf1  | 4200000000 | Leaf | provisioned
| cfg in-sync | NA         | NA      | 2    | 1  |
| 10.20.246.4 | r1   | NH-Leaf2  | 4200000000 | Leaf | provisioned
| cfg in-sync | NA         | NA      | 2    | 1  |
+-------------+------+-----------+------------+------+-------------+-------------+------------
+---------+------+----+

efa fabric setting show --name fs-1 | grep -i loopback

| Loopback Scheme            | uniform               |
| Loopback IP Range          | 172.31.254.0/24       |
| RouterID Loopback IP Range | 172.31.128.0/24       |
| VTEP Loopback IP Range     | 172.31.64.0/24        |
```

```
| Loopback Port Number       | 1                    |
| VTEP Loopback Port Number  | 2                    |
```

```
Rack1-Device1(config)# do show
running-config ip router-id
ip router-id 172.31.254.62
NH-Leaf1(config)# do show running-
config interface Loopback
interface Loopback 1
 ip address 172.31.254.62/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.254.17/32
 no shutdown
!
Rack1-Device1(config)#  do show
running-config overlay-gateway
overlay-gateway fs-1
 ip interface Loopback 2
 no map vni auto
 map vlan 24 vni 30300
 map vlan 25 vni 30301
 map vlan 26 vni 30302
 map bridge-domain 4095 vni 30213
 map bridge-domain 4096 vni 30212
 activate
!
Rack1-Device1(config)#do show
running-config router bgp address-
family ipv4 unicast
router bgp
 address-family ipv4 unicast
  network 172.31.254.17/32
  network 172.31.254.62/32
  maximum-paths 8
  graceful-restart
 !
```

```
Rack1-Device2(config)# do show
running-config ip router-id
ip router-id 172.31.254.100
Rack1-Device2(config)# do show
running-config interface Loopback
interface Loopback 1
 ip address 172.31.254.100/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.254.17/32
 no shutdown
!
Rack1-Device2(config)# do show
running-config overlay-gateway
overlay-gateway fs-1
 ip interface Loopback 2
 no map vni auto
 map vlan 24 vni 30300
 map vlan 25 vni 30301
 map vlan 26 vni 30302
 map bridge-domain 4095 vni 30213
 map bridge-domain 4096 vni 30212
 activate
!
Rack1-Device2(config)# do show
running-config router bgp address-
family ipv4 unicast
router bgp
 address-family ipv4 unicast
  network 172.31.254.17/32
  network 172.31.254.100/32
  maximum-paths 8
  graceful-restart
```

Allocate IP using Granular Scheme

You can choose granular IP allocation scheme for Router ID and VTEP Loopback.

About This Task

Follow this procedure to allocate an IP for Router ID and VTEP Loopback using granular scheme.

Procedure

To allocate an IP for Router ID and VTEP Loopback using granular scheme, run the following command:

```
efa fabric setting update --name <fabric-name> --loopback-scheme <uniform|granular>
efa fabric setting update --name <fabric-name> --router-id-loopback-ip-range <ip-range>
--vtep-loopback-ip-range <ip-range>
```

Example

```
efa fabric show --name fs-1

Fabric Name: fs-1, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-success,
```

```
Fabric Health: Green
+-------------+------+-----------+------------+------+-------------+-------------+-----------
+---------+------+----+
| IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE | DEVICE
STATE |  APP STATE  | CONFIG    | PENDING | VTLB | LB |
|             |      |           |            |      |             |
| GEN REASON| CONFIGS | ID   | ID |
+-------------+------+-----------+------------+------+-------------+-------------+-----------
+---------+------+----+
| 10.20.246.3 | r1   | NH-Leaf1  | 4200000001 | leaf | provisioned
| cfg in-sync | NA        | NA      | 2    | 1  |
| 10.20.246.4 | r1   | NH-Leaf2  | 4200000001 | leaf | provisioned
| cfg in-sync | NA        | NA      | 2    | 1  |
+-------------+------+-----------+------------+------+-------------+-------------+-----------
+---------+------+----+


efa fabric setting show --name fs-1 | grep -i loopback

| Loopback Scheme            | granular              |
| Loopback IP Range          | 172.31.254.0/24       |
| RouterID Loopback IP Range | 172.31.128.0/24       |
| VTEP Loopback IP Range     | 172.31.64.0/24        |
| Loopback Port Number       | 1                     |
| VTEP Loopback Port Number  | 2                     |
```

```
Rack1-Device1(config)# do show
running-config ip router-id
ip router-id 172.31.128.77
Rack1-Device1(config)# do show
running-config interface Loopback
interface Loopback 1
 ip address 172.31.128.77/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.96/32
 no shutdown
!
Rack1-Device1(config)# do show
running-config overlay-gateway
overlay-gateway fs-1
 ip interface Loopback 2
 no map vni auto
 map vlan 24 vni 30300
 map vlan 25 vni 30301
 map bridge-domain 4094 vni 30212
 activate
!
Rack1-Device1(config)# do show
running-config router bgp address-
family ipv4 unicast
router bgp
 address-family ipv4 unicast
  network 172.31.64.96/32
  network 172.31.128.77/32
  maximum-paths 8
  graceful-restart
 !
!
Rack1-Device1(config)#
```

```
Rack1-Device2(config)# do show
running-config ip router-id
ip router-id 172.31.128.171
Rack1-Device2(config)# do show
running-config interface Loopback
interface Loopback 1
 ip address 172.31.128.171/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.96/32
 no shutdown
!
Rack1-Device2(config)# do show
running-config overlay-gateway
overlay-gateway fs-1
 ip interface Loopback 2
 no map vni auto
 map vlan 24 vni 30300
 map vlan 25 vni 30301
 map bridge-domain 4094 vni 30212
 activate
!
Rack1-Device2(config)# do show
running-config router bgp address-
family ipv4 unicast
router bgp
 address-family ipv4 unicast
  network 172.31.64.96/32
  network 172.31.128.171/32
  maximum-paths 8
  graceful-restart
 !
!
Rack1-Device2(config)#
```

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-enable {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-enable` 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-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, 4.1.0*.

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.

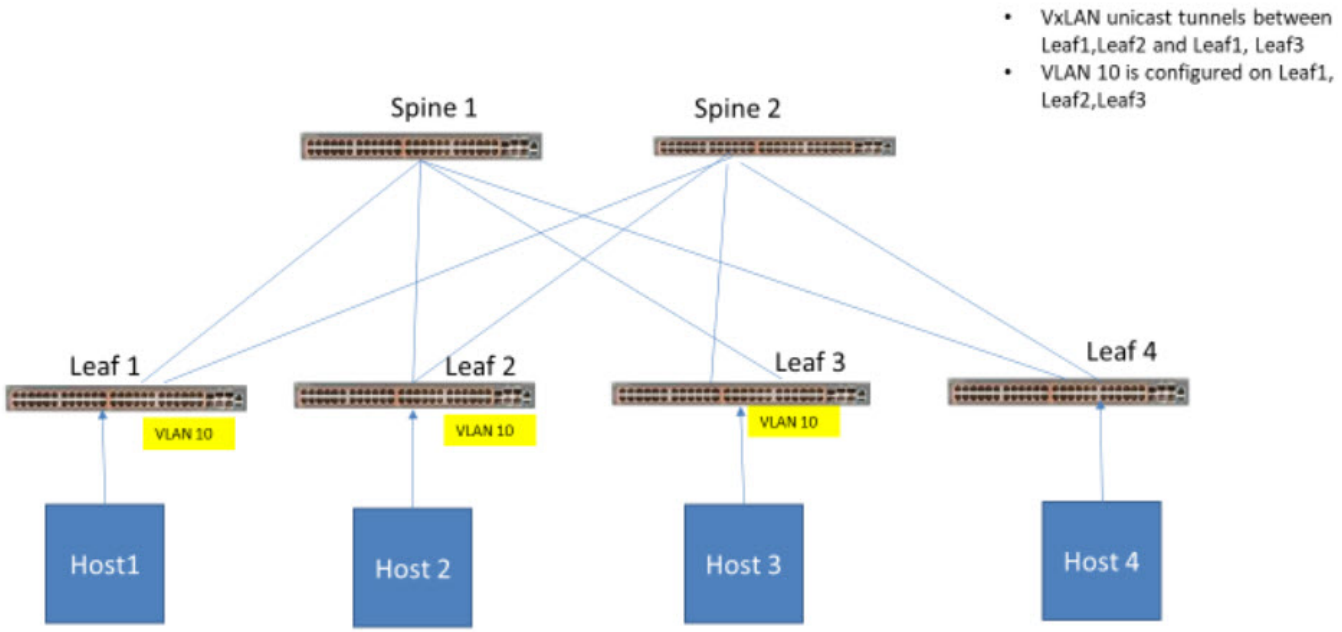

Figure 13: Clos topology with VxLAN unicast tunnels

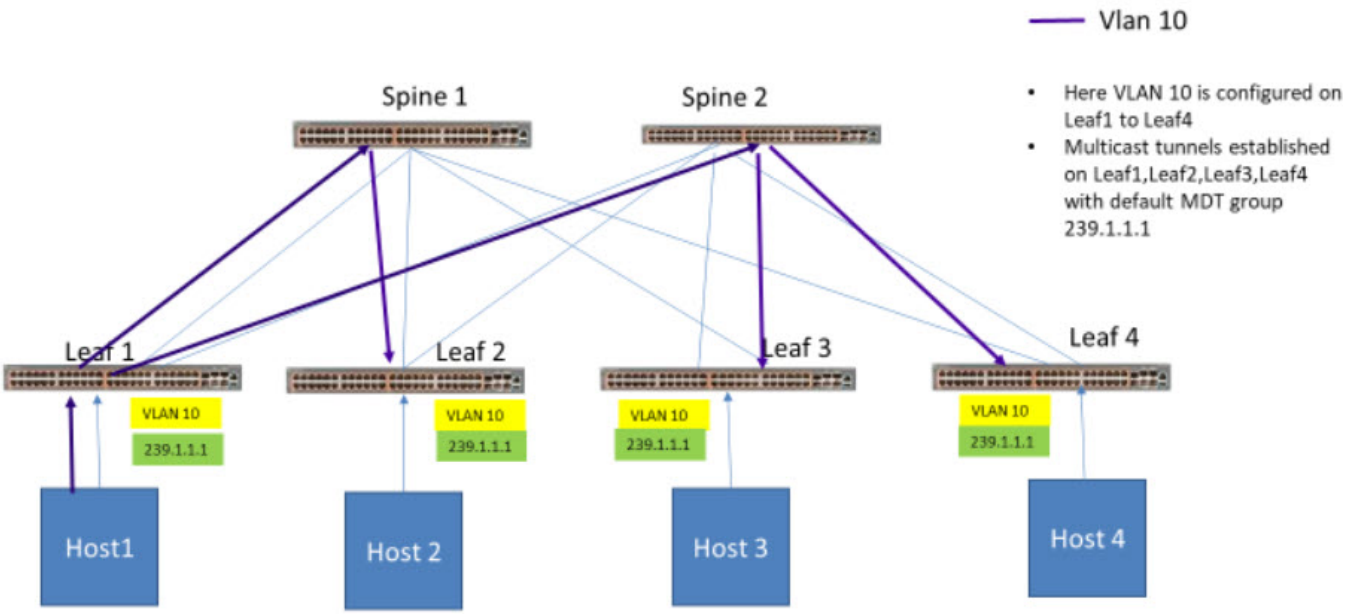

Figure 14: Clos topology with multicast tunnels

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 Configuration Guide, 4.1.0*.

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        |                               |
| Extreme 8720)                     |                                  |
+--------------------------------+-------------------------------+
| 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 <username> --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 193.

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
```

Link Add (LA) and Link Delete (LD) Enhancement

The enhancement ensures that when LA to LD or LD to LA events occur in sequence, devices transition to cfg-in-sync state immediately, without waiting for DRC or Fabric Configure. Stale configurations are cleaned up proactively during event handling.

It applies to all fabric topologies (with or without MCT):

- Non-Clos
- Clos Stage 3
- Clos Stage 5

LA and LD Behavioral Comparison: Before vs After Fix

The fix eliminates stale config artifacts, enables immediate cleanup, and ensures devices stay in-sync without waiting for fabric configure or DRC.

Link Add (LA)

Scenario	Before Fix	After Fix
Existing Link	ConfigType does not reset and remains stale in the following child tables: • `remote_neighbor_switch_config` • `interface_switch_config` • `member_cluster_config` • `cluster_members`	Intelligently resets ConfigType to NoneConfig for existing links.
New Link	ConfigType is set to CREATE_CONFIG.	ConfigType remains in CREATE_CONFIG.

Link Delete (LD)

Scenario	Before Fix	After Fix
Link added via LA, config never pushed	• Stale recipe remains until fabric configure runs. • On LD, system auto-cleans: ◦ Releases `used_ip_pair` and `used_ips` ◦ Deletes stale `DeleteConfig` entries from: ▪ `lldp_neighbor` ▪ `remote_neighbor_switch_config` ▪ `interface_switch_config` ▪ `member_cluster_config` ▪ `cluster_members`	• Device moves to `cfg-in-sync` immediately during LA event handler • No DRC or fabric configure required. • Recipes reset to NoneConfig if no pending config changes.
New LD	Recipes retained with DELETE_CONFIG.	Retains with DELETE_CONFIG for audit or trace.

Delete (LD) followed by Add (LA)

Before Fix	After Fix
• Device stuck in `cfg-refreshed` state until DRC or fabric configure runs. • Recipe unchanged, but device state (config) is out-of-sync.	• Device moves to `cfg-in-sync` immediately during the LA event handler. • No DRC or fabric configuration required. • Recipes reset to NoneConfig if no pending config changes.

Add (LA) followed by Delete (LD)

Before Fix	After Fix
• Device stays in `cfg-refreshed` state until fabric configure runs. • Recipe unchanged, but cleanup deferred.	• Device moves to `cfg-in-sync` immediately during LD event handler. • No fabric configure required. • Immediate cleanup performed: ◦ Releases `used_ip_pair` and `used_ips`. ◦ Purges DeleteConfig entries from all tables.

Interface Update (IU)

Before Fix	After Fix
IU is ignored if the interface is in DeleteConfig state because ConfigType block is processing.	IU is now processed even during the DeleteConfig state, which ensures that the interface state stays accurate without any loss across all events.

Existing Behavior: Dynamic Recipe Handling

Dynamic recipes are generated or updated during LinkAdd, LinkDelete, and InterfaceUpdate events. Each event triggers recipe regeneration with the appropriate ConfigType.

LinkAdd Event

If the link is new:

- A new recipe is created.
- Recipes for lldp_neighbors, interface_switch_configs, remote_neighbor_switch_configs, cluster_members, and mct_cluster_configs are prepared or updated.
- All associated recipes are marked with `ConfigType = CreateConfig`.
- Device and neighbor config gen reason set to LA.
- IP pairs are allocated for the new link (recorded as `used_ip_pair`)

LinkDelete Event

All relevant recipes for the deleted link or interface are blindly marked with `ConfigType = DeleteConfig`.

- Device and neighbor config gen reason set to LD.

InterfaceUpdate (IU) Event

Triggered only if:

- Interface speed changes
- Interface belongs to a multihomed device

InterfaceUpdates are applicable only if the Interface is part of the fabric and updates the recipes for `mct_cluster_configs` and `cluster_members`. The corresponding device config gen reason set to IU.

InterfaceAdd or InterfaceDelete Events

Applies to logical interfaces only.

- **InterfaceAdd:** Sets `ConfigType = CreateConfig`
- **InterfaceDelete:** Sets `ConfigType = NoneConfig` (or `UpdateConfig` if updating)
- Device and neighbor appstate marked with `cfg-refreshed`, includes `Config Gen Reason`.

.

Post-Event Processing

After all events are processed, DRC (Drift and Reconciliation) runs:

- Compares the switch configs (from inventory) with the fabric recipes.
- If there is no drift, device app state transitions to in sync, and all config gen reasons are cleared.

The following table differentiates between Config-Only and LLDP-Driven Config tables:

Config Only Table	LLDP Neighbor Driven Table
Bfd_profiles	Lldp_neighbors
Bgp_evpn_afs	Remote_neighbor_switch_config
Bgp_ip_afs	Interface_switch_config
Device_global_props	Mct_cluster_config
evpn	Cluster_members
Interface_vrf_mappings	Interface_vrf_mappings (Only for OS ONE)
Overlay_gateways	Ip_allocation_pools
Overlar_topology_intended	Ip_pair_allocation_pools
Router_bgp	Ip_prefix_list
Router_pims	Peer_groups
Switch_config	Phys_interfaces
Used_asn	Used_ip_pair
	Used_ips

DRC updates ConfigTypes of the following tables:

- Device_global_props
- Lldp_neighbors
- remote_neighbor_switch_config
- Mct_cluster_config
- Cluster_members

ConfigType Transitions & Bit-Field Logic Enhancement

In every event handler (for example, LA, LD, and IU), after the staging operation completes, the handler evaluate current and previous ConfigType state and apply smart transition rules using a new per-recipe bit-field to track the **modified-while-deleted** scenarios.

Bit-Field: Each recipe has a bit-field that tracks the config change while interface was in `DeleteConfig` state. This allows the system to distinguish between **clean delete** and **modified-while-deleted** scenarios.

ConfigType Transition Rules

- If `ConfigType = CreateConfig` AND previous state ≠ DeleteConfig, retain `ConfigType = CreateConfig`.
- If `ConfigType = CreateConfig` AND previous state = DeleteConfig AND bit-field = 0, transition to `ConfigType = NoneConfig`.
- If `ConfigType = CreateConfig` AND previous state = DeleteConfig AND bit-field > 0, transition to `ConfigType = UpdateConfig`.

Bit-Field Updates

Even if interface is in the `DeleteConfig` state, the bit-field updates are triggered by the following events:

- BGP IP neighbor changes
- Interface speed updates

Key Enhancements

- **Immediate Cleanup:** Stale recipes purge during LD and do not delay till DRC or fabric configure.
- **Bit-field Tracking:** Preserves config changes made while interface is in DeleteConfig.
- **ConfigType Intelligence:** Dynamic transitions prevent unnecessary drift and enable faster `cfg-in-sync` convergence.
- **No DRC Dependency:** Devices auto-sync during event handlers resulting in reduced latency and improved stability.

Use Cases Breakdown

Use this topic to learn about the flow of the event such as Link Add (LA), LD to LA (Delete followed by Add), and LA to LD (Add followed by Delete).

Use Case 1: Link Add (LA): New Link

- New link detected
- New recipe created
- Recipes updated:
  - `interface_switch_configs`
  - `remote_neighbor_switch_configs`
  - `cluster_members`
  - `mct_cluster_configs`
- All recipes marked with `ConfigType = CreateConfig`
- Device & neighbor `config_gen_reason` set to LA
- At the end of LA handler:
  - `ConfigType = CreateConfig`
  - `Bit-field = 0`

- Interface remains active and cannot be cleared
- Device stays `cfg-refreshed` until next event resolves drift.

Use Case 2: LD followed by LA (Delete to Add)

LD Phase

- All recipes marked with `ConfigType = DeleteConfig`
- `config_gen_reason = LD` on both devices

LA Phase

- Recipes regenerated with `ConfigType = CreateConfig`
- At the end of LA handler:
  - If prev state = LD and `bit-field = 0`, set `ConfigType = NoneConfig`.
  - Count remaining links:
    - If all links are NoneConfig, clear LD and LA from `config_gen_reason`.
    - If only BGPU remains. stay `cfg-refreshed`.
    - If all reasons cleared, move to `cfg-in-sy.c`.

Use Case 3: LA followed by LD (Add to Delete)

LA Phase

Recipes marked CreateConfig (as per Use Case 1).

LD Phase

- At the end of LD handler:
  - Remove entries from
    - ldp_neighbors
    - interface_switch_config
    - cluster_members
  - Release IPs:
    - Remove from `used_ip_pairs` or `used_ips`.
    - Return to `ip_pair_allocation_pool` or `ip_prefix_list`.
  - Set `remote_neighbor_switch_config, mct_cluster_config` to NoneConfig (if no drift)
  - Count remaining links:
    - If none in non-NoneConfig, clear LA or LD from `config_gen_reason`.
    - If all reasons cleared, move to `cfg-in-sync`.

Multi-Drift Use Case: Parent-Child (LD Followed by LA with BGP Update)

This use case covers a link deletion (LD) followed by a link addition (LA), while a BGP neighbor configuration change (for example, MD5 password update) is still pending.

The use case ensures correct handling of ConfigType, config_gen_reason, and device state transition to cfg-in-sync even with the overlapping config drifts.

Goal

- Correct handling of `ConfigType` and `config_gen_reason`
- Device state transitions to `cfg-in-sync` reliably
- No config loss during overlapping config drifts

Step-By-Step Flow

1. Link Delete (LD)
2. BGP IP Neighbor Update (BGPU) during LD while the link is in DELETE_CONFIG state
3. Link Add (LA)
   - Upon LA, the system re-activates link and regenerates recipes.
   - At the end of LA stage 4 of LA processing, all the affected recipes are set to `ConfigType = CREATE_CONFIG`.
   - At the end of LA handler:
     - If the link was previously in LD and the `bit-field = 0` (no pending config changes), the recipe's ConfigType is reset to NoneConfig.
     - If the bit-field > 0, the recipe remains in CREATE_CONFIG or switch to UPDATE_CONFIG per transition rules.
   - System checks if any other links on the device still have `ConfigType ≠ NoneConfig`.
     - If all links are clean, the system clears LD and LA from `config_gen_reason`.
     - If any other `config_gen_reasons` exists (for example, BGPU), device stays `cfg-refreshed` until resolved.
4. BGP Update Resolved (BGPU Cleared)
   - When the BGP neighbor config is reverted or applied successfully, BGPU event handler clears BGPU flag from `config_gen_reason`.
   - Once all `config_gen_reasons` are cleared, the device is moved to `cfg-in-sync`.

Key Points

- Bit-field usage:
  - Tracks if config change occurred while interface was in DELETE_CONFIG.
  - Stops BGP updates from disappearing during LD to LA transitions.
- ConfigType Flow:
  - DELETE_CONFIG to CREATE_CONFIG (on LA)
  - NoneConfig (if bit-field = 0 and no other pending configs) OR UPDATE_CONFIG (if bit-field > 0)
- Device Sync Rule: Only moves to `cfg-in-sync` when ALL `config_gen_reasons` are cleared including BGPU.
- No DRC or fabric configure needed: Everything happens during event handlers and is instant, atomic, without any external dependencies.

MCT Update

During MCT update or create event, `Mct_cluster_config` table does the following operations:

- If diff is found, move the ConfigType to CREATE (not UPDATE).
- Ensures cluster config gets created fresh, not just updated.

DRC vs Fabric Configure

The following table describes difference between DRC and configure fabric:

Parameter	DRC	Configure fabric
Config type	Only move lldp, remote ngbr, and global configs if the reconciliation is done	Move all configure to None type
No form of configs (Configs which need to be deleted)	No	Yes

Note

- Device syncs immediately. It does not wait for DRC or Fabric Configure to sync.
- BGP config changes during LD are NOT LOST.
- Handles multi-drift (link + BGP) flawlessly.

Non-Clos Single-Rack Fabric with Two Leaf Nodes and MCT Connections

Current Approach	New Approach
When a Leaf reboot or MCT link bounce occurs (LD followed by LA or LA followed by LD), both nodes enter `cfg-refreshed` state and wait for the deep discovery to clear the `config_gen_reasons` (LA, LD, IU) and then slowly sync back (move the device back to cfg-in-sync).	Upon the occurrence of the above scenarios, devices immediately clear `config_gen_reasons` during the respective event handlers (LA or LD) without waiting for the deep discovery. Both leaf nodes transition directly to `cfg-in-sync` once all pending config reasons are resolved.

Verification

- Leaf reboot: Verify that both nodes auto-sync (cfg-in-sync)
- MCT interface removed: Re-add and verify that both stay in-sync
- A new MCT link was added and then removed. Verify that no config drift
- Both leaves reboot together: Verify that they still auto-recover
- MCT links flaps multiple times. Verify that there is no config-gen backlog

Non-Clos Multi-Rack Fabric

Current Approach	New Approach
When BGP sessions between devices across different racks go down and up (LD followed by LA) or get added and removed (LA followed by LD), both devices auto-clear config-gen reasons, snap back to `cfg-in-sync` state without deep discovery.	Devices clear the config gen reason and move back to `cfg-in-sync` without deep-discovery.

3-Stage Clos Fabric Without MCT

In a 3-Stage Clos fabric without MCT, the topology consists of spine nodes, leaf nodes, and BGP peering sessions between spine and leaf.

Current Approach	New Approach
When any of the following occurs: • A Spine node reboots (link down followed by link up) • A BGP interface is removed and re-added (LD followed by LA) • A new BGP link is added and later removed (LA followed by LD) Both Spine and Leaf nodes enter `cfg-refreshed` state. They wait for Deep Discovery to clear `config_gen_reason` flags and transition back to `cfg-in-sync`.	With the new design: • Devices (Spine and Leaf) clear `config_gen_reason` and transition to `cfg-in-sync` during the event handler itself without requiring Deep Discovery. • No dependency on periodic reconciliation or manual intervention.

Example: Spine Node Reload

- Spine node goes down > all BGP links go down > LD events generated.
- Spine node comes back up > BGP sessions re-established > LA events received by fabric.
- During LA event handling:
  - Recipes updated, ConfigType adjusted per transition rules.
  - config_gen_reason cleared for both Spine and Leaf.
  - Devices immediately move to cfg-in-sync without needing any deep discovery.

3-Stage Clos Fabric With MCT

In a 3-Stage Clos fabric with MCT, the topology consists of one spine node, two leaf nodes, BGP peering sessions between spine and both leaf nodes, and MCT (Multi-Chassis Trunk) links between the two Leaf nodes.

Current Approach	New Approach
When any of the following occurs: • Spine node goes down and comes back up • BGP interface is removed and re-added between Spine and Leaf nodes (one or both) • New BGP link is added and later removed between Spine and Leaf All affected devices (Spine + both Leafs) enter `cfg-refreshed` state. They wait for Deep Discovery to clear config_gen_reason flags and transition back to `cfg-in-sync`. Additionally, for MCT links between the two Leaf nodes, behavior mirrors that of Non-Clos Single-Rack Fabric: • LD followed by LA or LA followed by LD sequences cause temporary cfg-refreshed state. • Recovery requires Deep Discovery.	With the new design: • All devices (Spine and both Leafs) clear `config_gen_reason` and transition to `cfg-in-sync` during the event handler without requiring Deep Discovery. • Applies to both BGP links (Spine-Leaf) and MCT links (Leaf-Leaf) • Eliminates dependency on periodic reconciliation or manual intervention.

5-Stage Clos Fabric Without MCT

In a 5-Stage Clos fabric without MCT, the topology consists of super spine nodes, spine nodes, leaf nodes, BGP peering sessions between spine and leaf, super spine and spine, and MCT (Multi-Chassis Trunk) links between the two Leaf nodes (even without MCT at higher tiers).

Current Approach	New Approach
When any of the following occurs: • SuperSpine node goes down and comes back up • BGP interface is removed and re-added between SuperSpine and Spine • New BGP link is added and later removed between SuperSpine and Spine SuperSpine and Spine nodes enter `cfg-refreshed` state. They wait for Deep Discovery to clear `config_gen_reason` and transition back to `cfg-in-sync`. Additionally, • Between Spine and Leaf: Behavior matches 3-Stage Clos (without MCT) and requires Deep Discovery after LD/LA cycles. • Between two Leaf nodes (MCT): Behavior matches Non-Clos Single-Rack fabric and requires Deep Discovery. –	Upon any BGP or MCT link up/down event, all the affected nodes (SuperSpine, Spine, Leaf) immediately clear config-gen reason and transition to `cfg-in-sync` without triggering deep discovery. With the new design: • All affected devices, SuperSpine, Spine, and Leaf clear `config_gen_reason` and transition to `cfg-in-sync` during the event handler without requiring Deep Discovery. • Applies uniformly across: ◦ SuperSpine - Spine BGP links ◦ Spine - Leaf BGP links ◦ Leaf - Leaf MCT links • Immediate cleanup of stale recipes, IP pairs, and config flags during LD or LA handlers. –

5-Stage Clos Fabric With MCT

All previous use cases (Clos Stage 3, Non-Clos, and so on) apply here with MCT enabled at Leaf layer.

What Fixed	What Not Fixed
The design fix for LA/LD and LD/LA sequences ensures that devices transition to `cfg-in-sync` during event handlers, without relying on Deep Discovery, regardless of whether links go down and up or up and down.	The following remain out of scope: • DRC pushing "Yes" (CreateConfig) / "No" (DeleteConfig) configs • DRC removing configs marked for DeleteConfig from switches • DRC Drift Handling: Setting `ConfigType = NoneConfig` if no drift exists across all fabric table • Deep Discovery Ignoring Events: Still processed if triggered externally • Config-Only Event Handling: Not optimized separately

Dynamic Recipe Generation and Update Flow

Dynamic recipes are dynamically generated or updated during LinkAdd, LinkDelete, and InterfaceUpdate events. Each event triggers recipe regeneration or ConfigType updates to reflect the new network state, ensuring config consistency without manual intervention.

LinkAdd Event

- If the link is new, new recipe is created.
- The following recipes are prepared or updated:
  - lldp_neighbors
  - interface_switch_configs
  - remote_neighbor_switch_configs
  - cluster_members
  - mct_cluster_configs
- All marked with `ConfigType = CreateConfig`.
- Device & neighbor `config_gen_reason = LA`.
- IP pairs allocated and stored in `used_ip_pair`.

LinkDelete Event

- All associated recipes marked with `ConfigType = DeleteConfig`.
- Device & neighbor `config_gen_reason = LD`.

InterfaceUpdate (IU) Event

- Applies only if:
  - Interface is part of the fabric
  - Speed changed
  - Belongs to a multihomed device
- Updates:
  - mct_cluster_configs
  - cluster_members
- Device config_gen_reason = IU.

InterfaceAdd or InterfaceDelete (Logical Interfaces Only)

- InterfaceAdd: `ConfigType = CreateConfig`
- InterfaceDelete: `ConfigType = UpdateConfig` or `NoneConfig`
- Device & neighbor:
  - App state changed to `cfg-refreshed`
  - Updated with respective `config_gen_reason`

Post-Event Processing: DRC Execution

- After all events are processed, DRC runs.
- DRC compares Inventory switch configs with Recipe-derived configs.

- If no drift, devices move to cfg-in-sync state and all the config gen reasons are cleared.
- If drift is detected, corrective actions are triggered (per existing logic).

ConfigType Transition Rules (Stage 5 Enhancement)

The system uses a bit-field to track config changes while the interface is in DeleteConfig state.

1. CreateConfig (Not Previously DeleteConfig)

   If the current ConfigType is CreateConfig and the previous ConfigType was NOT DeleteConfig, then retain CreateConfig as the active ConfigType.
2. CreateConfig (Previously DeleteConfig, BitField = 0)

   If the current ConfigType is CreateConfig and the previous ConfigType was DeleteConfig and the bit-field = 0, then transition the ConfigType to NoneConfig.
3. CreateConfig (Previously DeleteConfig, BitField > 0)

   If the current ConfigType is CreateConfig and the previous ConfigType was DeleteConfig and the bit-field > 0, then transition the ConfigType to UpdateConfig.

Interface IP Allocation Scenarios

In one scenario after a link delete (LD) event followed by fabric service restart and link add (LA), the physical interface IPs went empty from the phys_interface table. When new IPs were generated and allocated, the DRC pushed duplicate IPs. The similar issue is observed after a link add (LA) event. This issue caused IP collisions and prevented the fabric from stabilizing correctly.

Root Cause	Fix Implemented
• After fabric restart, the in-memory state of physical interface IPs was lost. • The system incorrectly assumed no prior IP allocation existed. • New IPs were allocated during LA. Later, DRC attempted to reconcile which pushed the old IPs again and created an IP conflict.	• Prevented physical interface IPs from being cleared after fabric restart. • Never assign conflicting IPs. If `phys_interface` IPs are missing or mismatched, fall back to `interface_switch_config`. • Ensured IPs persist across restarts, avoiding re-allocation during LA. • Eliminated duplicate IP assignments and ensured no more manual DRC conflicts.

Known Issues

- Does not recover from error states, only prevents entry into them.
- Does not handle heterogeneous PO speeds and fix upgrade path from pre-v4.0.2.
- The Link add event followed by the link delete event leaves VRU event in `config-gen-reaso`

While fixing (LD to LA scenario), nine real-world scenarios were analyzed to prevent IP collisions and ensure consistent, correct IP pair assignment across link add/delete cycles, even after restarts or cross-connects.

#1 Same Link: Both Phys_Interface Table IPs Present (LD Followed by LA)

Behavior	Example

If IPs already exist in phys_interface → just reassign the same IPs. Simple. Predictable. Stable.	• Existing Link
	`D1 (0/1): 10.10.10.0 ↔ D2 (0/2): 10.10.10.1`
	• Link deleted and then re-established (on same devices, same ports)
	Reuse the same pair 10.10.10.0 / 10.10.10.1 because the `phys_interface` still has
	`D1 (0/1): 10.10.10.0`
	`D2 (0/2): 10.10.10.1`

#2 Same Link: Both Phys_Interface Table IPs Empty (LD Followed by LA)

Behavior	Example

When there are no IPs in `phys_interface`, system falls back to `interface_switch_config` to find the last known good IPs for this exact link. The system reuses the IPs if it exists, and allocate a fresh pair of IPs if no IPs are found.	• Existing Link
	`D1 (0/1): 10.10.10.0 ↔ D2 (0/2): 10.10.10.1`
	• Link deleted and then re-established
	If the `phys_interface` is empty, check the `interface_switch_config` table. If found, reuse the iPs, else use a new pair.
	`D1 0/1 ↔ D2 0/2: 10.10.10.0 / 10.10.10.1`

#3 Same Link: One Phys_Interface IP Empty (LD Followed by LA)

Behavior	Example

If one side's IP is gone (empty in phys_interface), we derive the peer IP from the existing IP — and reuse the same pair.	• Existing Link
	`D1 (0/1): 10.10.10.0 ↔ D2 (0/2): 10.10.10.1`
	• New Link
	`D1 (0/1): 10.10.10.0`
	`D2 (0/2): empty`

#4 Cross-Connect on Same Device and Different Interface: Phys IPs Present

Behavior	Example

Even though both ends have IPs, but they are not in the same subnet, do not reuse, instead allocate fresh IP pair.	• Existing Links (D1 to D2)
	`0/1 ↔ 0/2 → IPs: 10.10.10.0 / 10.10.10.1`
	`0/3 ↔ 0/4 → IPs: 10.10.10.2 / 10.10.10.3`
	• New Link
	`0/1 ↔ 0/4 (same devices, different interfaces — cross-connect)`
	Note: Both interfaces have IPs, but 10.10.10.0 (0/1) and 10.10.10.3 (0/4) are not in same subnet. So this is an invalid pair for this link. So, assign new IPs: 10.10.10.4 / 10.10.10.5.

#5 Cross-Connect on Same Device and Different Interface: One Phys_Interface IP Empty

Behavior	Example	Solution

When one side has an IP but the other does not, the system tries to derive peer IP from the existing one. But, the system picks up the wrong or matched IP which causes an IP conflict.	• Existing Links (same device D1
	`0/1 ↔ D2: 10.10.10.0 / 10.10.10.1`
	`0/3 ↔ D2: 10.10.10.2 / 10.10.10.3`
	• New Link
	`0/1 ↔ 0/4 (cross-connect — D1's 0/1 to D2's 0/4)`
	For example, when `phys_interface` for 0/1 has IP (10.10.10.0), but 0/4 is empty, system derives peer IP as 10.10.10.1 (from 0/1's pair), But 10.10.10.1 belongs to 0/2, not 0/4, which is a wrong pair and creates a conflict.
		Before assigning derived IP, we cross-check with `interface_switch_config`:
		• Does the derived peer IP match what's recorded for that interface?
		• Only if it matches, reuse the pair.
		• Else, allocate new IPs: 10.10.10.4 / 10.10.10.5

#6 Cross-Connect on Same Device and Interface: Both Phys_Interface IPs Empty

Behavior	Example

Fallback to interface_switch_config → but if derived IPs aren't in same subnet → allocate new IP pair.	• Existing Links (same device D1)
	`0/1 ↔ D2: 10.10.10.0 / 10.10.10.1`
	`0/3 ↔ D2: 10.10.10.2 / 10.10.10.3`
	• New Link
	`0/1 ↔ 0/4 (cross-connection on same device — D1's 0/1 to D2's 0/4)`
	
	Note: When the `phys_interface` IPs are empty, the system checks `interface_switch_config` and finds 10.10.10.2/3. But those belong to 0/3-0/4, not 0/1-0/4, which results in a subnet mismatch. So, assign fresh IPs: 10.10.10.4 / 10.10.10.5.

#7 Cross-Connect on Different Device and Interface): Phys_Interface IPs Present

Behavior	Example

Same as Scenario #4. since IPs exist but are not in the same subnet, system allocates new IP pair to avoid conflicts.	• Existing Links
	
	(All on separate devices without overlap)
	`- D1 (0/1) ↔ D2 (0/2) → IPs: 10.10.10.0 / 10.10.10.1`
	`D3 (0/3) ↔ D4 (0/4) → IPs: 10.10.10.2 / 10.10.10.3`
	• New Link
	`D1 (0/1) ↔ D4 (0/4) - cross-device, cross-interface`
	
	Note: The `phys_interface` has IPs, but 10.10.10.0 (D1) and 10.10.10.3 (D4) are not in same subnet, which creates an invalid pair. So, assign fresh IPs: 10.10.10.4 / 10.10.10.5.

#8 Cross-Connect on Different Device and Interface: One Phys_Interface IP Empty

Behavior	Example	Solution

Same as Scenario #5, but with inter-device cross-connect.	• Existing links
	`- D1: 10.10.10.0 (0/1) ↔ D2: 10.10.10.1 (0/2) D3: 10.10.10.2 (0/3) ↔ D4: 10.10.10.3 (0/4)`
	• New cross-connect link
	`D1 (0/1) ↔ D4 (0/4) - different devices, different ports`
		Added validation check:
		• Compare derived peer IP with what's stored in `interface_switch_config`.
		• Only reuse if they match (which only happens in Scenario #3 - same link, one IP missing). Otherwise, allocate fresh IPs: 10.10.10.4 / 5.

#9 Cross-Connect on Different Device and Interface: Both Phys_Interface IPs Empty

Behavior	Example

Same as Scenario #6: Fallback to `interface_switch_config`, but if derived IPs are not in subnet, allocate **new IP pair**.	• Existing links
	`D1: 10.10.10.0 (0/1) and D2: 10.10.10.1 (0/2) D3: 10.10.10.2 (0/3) and D4: 10.10.10.3 (0/4)`
	• New cross-connect link
	`D1 (0/1) ↔ D4 (0/4) - different devices, different interfaces`
	
	When the `phys_interface` IPs are empty, system checks `interface_switch_config` and finds 10.10.10.2/3. But, those belong to D3-D4, not D1-D4, which results in subnet mismatch. So, allocate fresh IPs: 10.10.10.4 / 5.
	
	Note: This fix prevents the system from entering an error state. It doesn't repair broken setups.

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 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 output of **`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 |
|              |       |             |           |      |              |             | GEN REASON |
```

```
CONFIGS  | ID   | 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 the `loopback-ip-range`, `p2p-link-range`, and `mctlink-ip-range` fabric settings. Within the same network, ensure that the prefix mask length is lower than the configured prefix mask length.

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>`
  - `--bfd-session-type`

  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 --allow-as-in --mtu --ip-mtu --rack-underlay-ebgp-peer-group --rack-overlay-ebgp-peer-group --lacp-timeout --control-vlan --control-ve --vni-auto-map --backup-routing-enable --backup-routing-ipv4-range --backup-routing-ipv6-range --overlay-gateway-broadcast-local-bias --loopback-scheme --router-id-loopback-ip-range <ip-range> --vtep-loopback-ip-range <ip-range> --reserve-first-4k-vni-for-vlan	--p2p-link-range --loopback-ip-range --bfd-enable --bfd-tx --bfd-rx --bfd-multiplier --bfd-session-type --mctlink-ip-range --single-rack-deployment --md5-password-enable --md5-password --bgp-dynamic-peer-listen-limit --bgp-multihop --rack-mct-scheme --rack-mct-ports --rack-ld-mct-ports --max-paths --mct-peer-keep-alive-enable --mct-port-channel	--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 --allow-as-in --mtu --ip-mtu --leaf-peer-group --spine-peer-group --super-spine-peer-group --lacp-timeout --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 --overlay-gateway-broadcast-local-bias --loopback-scheme --router-id-loopback-ip-range <ip-range>	--p2p-link-range --loopback-ip-range --bfd-enable --bfd-tx --bfd-rx --bfd-multiplier --bfd-session-type --mctlink-ip-range --md5-password-enable --md5-password --bgp-dynamic-peer-listen-limit --max-paths --mct-peer-keep-alive-enable --mct-port-channel	--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 --rack-mct-scheme --rack-mct-ports --rack-ld-mct-ports

Cannot be Updated	Can be Updated	Invalid Settings for Clos Fabric
`--vtep-loopback-ip-range <ip-range>` `--reserve-first-4k-vni-for-vlan`		

MD5 Password Encoding in Fabric Settings

Beginning with XCO 4.1.0, XCO encodes the BGP MD5 password before storing it in the XCO database. The plain-text password you enter is never written to the database, displayed in command output, or written to logs. As a result, the MD5 password always appears in an encoded or hardware-encrypted form wherever it is shown. For example, in the output of **`efa fabric settings show advanced`** and **`efa fabric show-running-config`**.

Note
An encoded password in command output or logs is expected behavior. It does not indicate that the password is corrupted or misconfigured.

Password Encoding Prefixes

XCO uses a prefix on the stored value to indicate how the MD5 password is protected. Two prefixes can appear in command output:

Prefix	Meaning	Origin
$#$	Software-encoded by XCO (base64).	Applied by XCO when you set the password with the `--md5-password` option.
9	Hardware-encrypted by the SLX device.	Returned by the SLX device. XCO stores and displays the value as-is.

A value shown without either prefix is plain text. Plain text is accepted only as input when you set a password; it is never stored or displayed.

Setting MD5 Password

Set the fabric BGP MD5 password with the `--md5-password` option of **`fabric settings update`**:

```
efa fabric settings update --name myfabric
--md5-password "myBgpSecret"
```

XCO encodes the value with the $#$ prefix and stores only the encoded form. The plain-text value you entered (in this example, `myBgpSecret`) is not persisted.

Viewing Stored Password

The **efa fabric show-running-config** and **efa fabric settings show advanced** commands display the stored (encoded or hardware-encrypted) form of the MD5 password, not the plain text you entered:

```
efa fabric show-running-config --name myfabric

# Output shows:
--md5-password '$#$bXlCZ3BTZWNyZXQ='
```

The `$#$...` value is the software-encoded form of the password. If the password originated from an SLX device, the value is displayed with a 9 prefix instead.

Reusing Show-Running-Config Output

You can copy **efa fabric show-running-config** output, including the encoded $#$... password and re-apply it to the same fabric or use it to recreate fabric settings on another system. XCO detects the $#$ prefix and stores the value as-is, without encoding it a second time. Re-applying an already-encoded value therefore does not double-encode the password and produces no unintended change.

Hardware-encrypted 9 values are handled the same way: XCO stores and re-applies them transparently.

Note
Because XCO accepts the encoded value directly, you do not need to know or re-enter the original plain-text password to reproduce a fabric's settings from **show-running-config** output.

Security

- The plain-text MD5 password is never stored in the XCO database.
- The plain-text MD5 password is never written to command output or logs.
- After a password is set, it appears only in software-encoded ($#$) or hardware-encrypted (9) form.

Troubleshooting

If you see a `$#$...` or `$9$...` value where an MD5 password appears in command output or logs, this is the expected encoded or hardware-encrypted form and not a sign of corruption or misconfiguration. To change the password, run the **efa fabric settings update** command again with a new plain-text value; XCO re-encodes and stores the new value.

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.

Update max-path on an Active Fabric

You can update the max-path on an active fabric.

About This Task

Follow this procedure to update the max-path on an active fabric.

DRC Behavior Limitation

After updating fabric settings, running DRC will not apply changes to all end-devices. Instead, use the **`efa fabric configure`** command to push updates. However, max-path setting updates are an exception: DRC will unexpectedly apply them to all end-devices even when the **`efa fabric configure`** command is not run.

Procedure

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

```
efa fabric setting update --name fab2 --max-paths 12
fab2 Fabric Update Successful
--- Time Elapsed: 16.864156868s ---
```

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

```
efa fabric setting show --name fab2 --advanced
+---------------------------------+-------------------------+
|               NAME              |          VALUE          |
+---------------------------------+-------------------------+
| Fabric Name                     | fab2                    |
+---------------------------------+-------------------------+
| MaxPaths                        | 12                      |
+---------------------------------+-------------------------+

efa fabric show --name fab2

Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status: settings-updated,
Fabric Health: Red

Updated Fabric Settings: BGP-MAXPATH

+--------------+-------+----------+------------+------+--------------+-------------
+------------------+----------------+---------+-------+
|   IP ADDRESS   | RACK  | HOST NAME |    ASN
| ROLE | DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-------+----------+------------+------+--------------+-------------
+------------------+----------------+---------+-------+
| 10.64.204.45 | rack1 | SLX       | 4200000000 |
leaf | provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.204.46 | rack1 | SLX       | 4200000000 |
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, MCT-PORTS - Rack Mct Ports, LD-MCT-PORTS - Rack Low Density Mct Ports
BGP-MAXPATH - Bgp Max Path, MCT-PK-ENABLE - MCT Peer keep Alive Enable
```

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

```
efa fabric configure --name fab2
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: 1m8.620487895s --4
```

4. Run the **efa fabric setting show** command.

```
efa fabric setting show --name fab2 --advanced
+--------------------------------+-------------------------+
|              NAME              |          VALUE          |
+--------------------------------+-------------------------+
| Fabric Name                    | fab2                    |
+--------------------------------+-------------------------+
| MaxPaths                       | 12                      |
+--------------------------------+-------------------------+
```

5. Complete the following configuration on SLX nodes.

```
#10.64.196.53
router bgp
..
 address-family ipv4 unicast
  ….
  maximum-paths 12
  graceful-restart
 !
!
```

```
#10.64.196.55
router bgp
..
 address-family ipv4 unicast
  ….
  maximum-paths 12
  graceful-restart
 !
!
```

Enable or Disable MCT Peer Keep Alive on an Active or Inactive Fabric

You can enable or disable MCT peer keep alive on an active or inactive fabric.

About This Task

Follow this procedure to enable or disable MCT peer keep alive configuration on an active or inactive fabric.

Note
For fabrics running versions prior to XCO 3.8.5, upgrading to XCO 3.8.5 or later will automatically set the fabric setting value to "Yes".

Procedure

1. Create a fabric with default peer keep alive fabric setting.

```
efa fabric create --name fab2 --type non-clos
efa fabric show --name fab2
Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
created, Fabric Health: Red
+--------------+------+----------+------------+------+----------------+-----------
+------------------+-----------------------------------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE |  DEVICE STATE   | APP
STATE | CONFIG GEN REASON |           PENDING CONFIGS            | VTLB ID | LB ID |
+--------------+------+----------+------------+------+----------------+-----------
+------------------+-----------------------------------------+---------+-------+

efa fabric setting show --name fab2 -advanced
+--------------------------------+-------------------------+
|              NAME              |          VALUE          |
```

```
+--------------------------------+------------------------+
| Fabric Name                    | fab2                   |
+--------------------------------+------------------------+
…..
+--------------------------------+------------------------+
| MCT Peer Keep Alive            | Yes                    |
| Enable                         |                        |
+--------------------------------+------------------------+
```

2. Disable peer keep alive fabric setting on an in-active fabric.

```
efa fabric setting update --name fab2 --mct-peer-keep-alive-enable  No
fab2 Fabric Update Successful
--- Time Elapsed: 16.864156868s ---


efa fabric setting show --name fab2 -advanced
+--------------------------------+------------------------+
|              NAME              |         VALUE          |
+--------------------------------+------------------------+
| Fabric Name                    | fab2                   |
+--------------------------------+------------------------+
…..
+--------------------------------+------------------------+
| MCT Peer Keep Alive           | No                     |
| Enable                         |                        |
+--------------------------------+------------------------+
```

3. Add multiple devices to the fabric.

```
efa fabric device add-bulk --name fab2 --rack r1
           --ip 10.64.196.53-54 --username admin --password password

efa fabric show --name fab2
Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
created, Fabric Health: Red
+--------------+------+-----------+------------+------+-----------------+-----------
+-------------------+----------------------------------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE |  DEVICE STATE   | APP
STATE | CONFIG GEN REASON |           PENDING CONFIGS            | VTLB ID | LB ID |
+--------------+------+-----------+------------+------+-----------------+-----------
+-------------------+----------------------------------------+---------+-------+
| 10.64.196.54 | r1   | SLX       | 4200000000 | leaf | not provisioned | cfg
ready | DA                | SYSP-C,MCT-C,MCT-PA,BGP-C,EVPN-C,O-C | 2       | 1     |
| 10.64.196.53 | r1   | SLX       | 4200000000 | leaf | not provisioned | cfg
ready | DA                | SYSP-C,MCT-C,MCT-PA,BGP-C,EVPN-C,O-C | 2       | 1     |
+--------------+------+-----------+----------



efa fabric setting show --name fab2 -advanced
+--------------------------------+------------------------+
|              NAME              |         VALUE          |
+--------------------------------+------------------------+
| Fabric Name                    | fab2                   |
+--------------------------------+------------------------+
…..
+--------------------------------+------------------------+
| MCT Peer Keep Alive            | No                     |
| Enable                         |                        |
+--------------------------------+------------------------+
```

4. Configure the fabric.

```
efa fabric configure --name fab2
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: 1m8.620487895s ---

efa fabric show --name fab2

Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
configure-success, Fabric Health: Black
+--------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE | DEVICE
STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.196.54 | r1   | SLX       | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.196.53 | r1   | SLX       | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

5. Complete the following configuration on SLX devices.

```
#10.64.196.53
cluster fab2-cluster-1
 peer 10.20.20.2
 peer-interface Port-channel 63
member vlan all
 member bridge-domain all
!
!
```

```
#10.64.196.54
cluster fab2-cluster-1
 peer 10.20.20.3
 peer-interface Port-channel 63
member vlan all
 member bridge-domain all
!
```

6. Enable peer keep-alive fabric setting on an active fabric.

```
efa fabric setting update --name fab2 --mct-peer-keep-alive-enable  Yes
fab2 Fabric Update Successful
--- Time Elapsed: 16.864156868s ---

efa fabric show --name fab2

Fabric Name: fabs, Fabric Description: , Fabric Type: non-clos, Fabric Status:
settings-updated, Fabric Health: Black
Updated Fabric Settings: MCT-PKA-ENABLE
+--------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE | DEVICE
STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.196.53 | r1   | SLX53     | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.196.54 | r1   | SLX54     | 4200000000 | 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, MCT-PORTS
- Rack Mct Ports, LD-MCT-PORTS - Rack Low Density Mct Ports
BGP-MAXPATH - Bgp Max Path, MCT-PK-ENABLE - MCT Peer keep Alive Enable

efa fabric setting show --name fab2 -advanced
+--------------------------------+-------------------------+
|              NAME              |          VALUE          |
+--------------------------------+-------------------------+
```

```
| Fabric Name                    | fab2                    |
+--------------------------------+-------------------------+
…..
+--------------------------------+-------------------------+
| MCT Peer Keep Alive            | yes                     |
| Enable                         |                         |
+--------------------------------+-------------------------+
```

7. Configure the fabric.

```
efa fabric configure --name fab2
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: 1m8.620487895s ---

efa fabric show --name fab2

Fabric Name: fab2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
configure-success, Fabric Health: Black
+--------------+------+-----------+------------+------+-------------+-------------
+------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE | DEVICE
STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------+-------------+-------------
+------------------+-----------------+---------+-------+
| 10.64.196.54 | r1   | SLX       | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.196.53 | r1   | SLX       | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+------+-----------+------------+------+-------------+-------------
+------------------+-----------------+---------+-------+
```

8. Complete the following configuration on SLX nodes.

```
#10.64.196.53
cluster fab2-cluster-1
 peer 10.20.20.2
 peer-interface Port-channel 63
 peer-keepalive
  auto
 !
 member vlan all
 member bridge-domain all
!!
```

```
#10.64.196.54
cluster fab2-cluster-1
 peer 10.20.20.3
 peer-interface Port-channel 63
 peer-keepalive
  auto
 !
 member vlan all
 member bridge-domain all
!
```

9. Display the fabric health with keep alive enabled and disabled.

Note

The operational state of keepalive can show as false even if keepalive has been configured. This is because the state is determined by both port-channel and keepalive settings. For instance, if the port-channel is down, the health status will be reported as false. This example demonstrates why the keepalive health might display as false, ensuring this condition is visible in the command output.

Additionally, when the KEEPALIVE fabric option is set to either Yes or No, if the overall cluster health status is green, the show command does not display child attributes. For this reason, output for the false state is included, so users can understand and compare differences in the show output under various conditions.

- When the peer keep alive is enabled.

```
efa fabric health show --name fabs --detail
================================================================================
Fabric Name                              :  fabs
Fabric Type                              :  non-clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
--------------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                        :  10.64.196.54 [Leaf]
  Device Health                          :  Black
  Configuration State Health             :  Black
    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
 Device IP [Role]                        :  10.64.196.53 [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
================================================================================
```

- When the peer keep alive is disabled.

```
efa fabric health show --name fabs --detail
================================================================================
========
Fabric Name                              :  fabs
Fabric Type                              :  non-clos
Fabric Health                            :  Black
 Fabric Status                           :  configure-success
 Fabric Level Physical Topology Health   :  Green
--------------------------------------------------------------------------------
--------
Fabric Device Health

 Device IP [Role]                        :  10.64.196.54 [Leaf]
  Device Health                          :  Black
  Configuration State Health             :  Black
    Dev State                            :  provisioned
    App State                            :  cfg in-sync
  Operational State Health               :  Black
    Cluster Health                       :  Black
      Operational State                  :  true
      Peer Operational State             :  false
      Physical Topology Device Health    :  Green
 Device IP [Role]                        :  10.64.196.53 [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
    Physical Topology Device Health     :  Green
================================================================================
========
```

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 BGPLL - BGP Listen Limit MD5 - BGP MD5 Password ASN - ASN Update
Overlay gateway	Cfg-refresh	OD/OU - Overlay Gateway Delete or Update
Interface (Phy/ Lo/Po)	Cfg-refresh	IA/ID/IU - Interface Add, Delete or Update POU - Port Channel Update
LLDP	Cfg-refresh	LA/LD - Link Add or Delete
Router Pim	Cfg-refresh	PC/PD/PU - RouterPim Create/Delete/Update
Prefix List	Cfg-refresh	PLC/PLD/PLU - IPPrefixList Create/Delete/Update
System	Cfg-refresh	SYS-System Properties Update (Hostname, MTU, IP MTU)
Device	Cfg-refresh	DD-Dependent Device Update, DA/DR - Device Add or ReAdd

Brownfield Fabric Service Overview

The Brownfield Fabric Service facilitates migration of existing fabric deployments to newer XCO versions, allowing legacy software to coexist with new software. Key capabilities include:

- **Migration from older XCO versions:** upgrade to newer XCO versions, even on a different server.
- **Migration of XCO from TPVM to external server:** transition XCO from TPVM to an external server.
- **Partial or full configuration migration:** migrate fabric configurations made via SLX CLI or out-of-band methods, with conflict checks.
- **Validation and error handling:** learns and validates device configurations, performs validation checks, and reports errors for any deviations.

Note
Brownfield deployments are not supported for Tenant Services.

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.

Use Case	Valid	Invalid
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.	
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.
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.

Preserve Retain Route Target All on Boarder Leaf Devices

XCO preserves the `retain route-target all` configuration on boarder leaf devices in any config sync operations (for example, any sync refresh from XCO to IP fabric).

In XCO 3.5.0, the `retain route-target all` configuration is applied under "l2evpn address-family" on the spine and super-spine devices, but not on the leaf and border-leaf devices. If you manually configure `retain route-target all` as out-of-band on the border-leaf or leaf devices, XCO detect this as a drift and reconcile.

Starting from XCO 3.6.0, for Data Center Interconnect (DCI) scenarios, ensure that the `retain route-target all` configuration is present on the border-leaf devices. Therefore, XCO will no longer manage the `retain route-target all` configuration, specifically, on the border-leaf devices. This means that XCO will neither configure nor detect drift and reconcile for any `retain route-target all` configuration on the border-leaf devices, except for non-Clos multi rack fabric.

The following table describes the `retain route-target all` configuration for a non-Clos and Clos fabric:

Fabric Type		Device Type	Description
Non-Clos	Singel Rack	Border Leaf	XCO does not push 'retain route-target all' on leaf or border-leaf devices, so it disowns the 'retain route-target all' configuration for these devices.
	Multi Rack	Border Leaf or Leaf	XCO preservers the 'retain route-target all' configuration, configures it as part of the fabric configuration, and detects and reconciles any drift during DRC.
Clos (3-Stage or 5-Stage)		Border Leaf	XCO disowns the 'retain route-target all' configuration, neither configuring it nor identifying and reconciling drift as part of DRC.

Example Configuration for Non-Clos Single Rack Fabric

The following is an example configuration for **non-Clos single rack fabric**:

```
efa fabric show --name common-fabric
Fabric Name: common-fabric, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-
success, Fabric Health: Green
+--------------+-------+-----------+------------+------------+--------------+-------------
+-------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK  | HOST NAME |    ASN     |    ROLE
| DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-------+-----------+------------+------------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.208.22 | r1    | SW-22     | 4200065535 | BorderLeaf
| provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.208.23 | r1    | SW-23     | 4200065535 | BorderLeaf
| provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+-------+-----------+------------+------------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

The following examples show the `retain route-target all` configuration behavior in XCO 3.5.0 and in XCO 3.6.0 and later in a **non-Clos single rack fabric**:

- **XCO 3.5.0 Behavior**

XCO CONFIGURED ROUTER BGP

```
SLX:
SW-22# show runn router bgp
router bgp
 local-as 4200065535
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200065535
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.211/32
  network 172.31.254.248/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

User created OOB configuration

```
SW-22# conf
Entering configuration mode
terminal
SW-22(config)# router bgp
SW-22(config-bgp-router)# address-
family l2vpn evpn
SW-22(config-bgp-evpn)# retain
route-target all
SW-22(config-bgp-evpn)# ex
SW-22(config-bgp-router)# ex
SW-22(config)# ex
```

```
efa inventory device update -ip 10.64.208.22
efa fabric show

Fabric Name: common-fabric, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-
success, Fabric Health: Green
+--------------+-------------+-----------+------------+------------+--------------+---------------
+------------------+-----------------+---------+-------+
|  IP ADDRESS  |    RACK     | HOST NAME |    ASN     |    ROLE    | DEVICE STATE |   APP
STATE   | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-------------+-----------+------------+------------+--------------+---------------
+------------------+-----------------+---------+-------+
| 10.64.208.22 | r1          | SW-22     | 4200065535 | BorderLeaf | provisioned  | cfg-refreshed
| BGPU             |                 | 2       | 1     |
| 10.64.208.23 | r1          | SW-23     | 4200065535 | BorderLeaf | provisioned  | cfg in-
sync   | NA               | NA              | 2       | 1     |
+--------------+-------------+-----------+------------+------------+--------------+---------------
+------------------+-----------------+---------+-------+


efa fabric configure --name common-fabric
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: 30.549013784s ---
```

SW-22# show runn router bgp
router bgp
 local-as 4200065535
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as 4200065535
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.211/32
  network 172.31.254.248/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
SW-22#
	The “retain route-target all” configuration gets removed from the border-leaf during “efa fabric configure” or during DRC.

- **Behavior in XCO 3.6.0 and later**

XCO CONFIGURED ROUTER BGP:
SLX:
SW-22# show runn router bgp
router bgp
 local-as 4200065535
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as 4200065535
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.211/32
  network 172.31.254.248/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
	OOB CREATED:
	SW-22# conf
	Entering configuration mode terminal
	SW-22(config)# router bgp
	SW-22(config-bgp-router)# address-family l2vpn evpn
	SW-22(config-bgp-evpn)# retain route-target all
	SW-22(config-bgp-evpn)# no graceful-restart
	%Warning: Please clear the neighbor session for the parameter change to take effect!
	SW-22(config-bgp-evpn)# ex
	SW-22(config-bgp-router)# ex
	SW-22(config)# ex

```
efa inventory device update -ip 10.64.208.22
efa fabric show -name common-fabric

Fabric Name: common-fabric, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-
success, Fabric Health: Green
+--------------+-------------+-----------+------------+------------+--------------+---------------
+-------------------+-----------------+---------+-------+
|  IP ADDRESS  |    RACK     | HOST NAME |    ASN     |    ROLE    | DEVICE STATE |   APP
STATE   | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-------------+-----------+------------+------------+--------------+---------------
+-------------------+-----------------+---------+-------+
| 10.64.208.22 | r1          | SW-22     | 4200065535 | BorderLeaf | provisioned  | cfg-in-
sync   | NA                | NA              | 2       | 1     |
```

```
| 10.64.208.23 | r1          | SW-23     | 4200065535 | BorderLeaf | provisioned  | cfg in-
sync   | NA               | NA              | 2       | 1     |
+--------------+-------------+-----------+------------+------------+--------------+--------------
+------------------+----------------+---------+-------+


efa fabric configure --name common-fabric
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: 30.549013784s ---
```

```
SW-22# show runn router bgp
router bgp
 local-as 4200065535
 capability as4-enable
 fast-external-fallover
 neighbor 10.20.20.5 remote-as
4200065535
 neighbor 10.20.20.5 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.211/32
  network 172.31.254.248/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
  retain route-target all
 !
!
SW-22#
```

The “retain route-target all” configuration doesn’t get removed from the border-leaf during “efa fabric configure”.

- **DRC**

```
efa inventory drift-reconcile execute --ip 10.64.208.22 --reconcile
+--------------+-----------+--------------------------------------+---------+--------+
|  IP ADDRESS  | RECONCILE |                 UUID                 | STATUS  | REASON |
+--------------+-----------+--------------------------------------+---------+--------+
| 10.64.208.22 | Yes       | 580570d5-11f5-4c51-be6c-c04346181081 | Success |        |
+--------------+-----------+--------------------------------------+---------+--------+
Drift and Reconcile Execute
Execute the CLI to get details : efa inventory drift-reconcile detail --uuid
580570d5-11f5-4c51-be6c-c04346181081
--- Time Elapsed: 11.348279ms ---

efa inventory drift-reconcile detail --uuid 580570d5-11f5-4c51-be6c-c04346181081
+--------------------------------+--------------------------------------+
|              NAME              |                VALUE                 |
+--------------------------------+--------------------------------------+
| UUID                           | 580570d5-11f5-4c51-be6c-c04346181081 |
+--------------------------------+--------------------------------------+
| Device IP                      | 10.64.208.22                         |
+--------------------------------+--------------------------------------+
| Status                         | success                              |
+--------------------------------+--------------------------------------+
| Execution Reason               | manual                               |
```

```
+--------------------------------+---------------------------------------+
| operation                      | drift-and-reconcile                   |
+--------------------------------+---------------------------------------+
| Inventory Status               | inventory-dr-success                  |
+--------------------------------+---------------------------------------+
| Is Inventory config Refreshed  | false                                 |
+--------------------------------+---------------------------------------+
| Inventory Duration             | 19.991554ms                           |
+--------------------------------+---------------------------------------+
| Fabric Status                  | fabric-dr-success                     |
+--------------------------------+---------------------------------------+
| Is Fabric config Refreshed     | false                                 |
+--------------------------------+---------------------------------------+
| Fabric Duration                | 33.609071659s                         |
+--------------------------------+---------------------------------------+
| Policy Status                  | policy-dr-success                     |
+--------------------------------+---------------------------------------+
| Is Policy config Refreshed     | false                                 |
+--------------------------------+---------------------------------------+
| Policy Duration                | 104.505439ms                          |
+--------------------------------+---------------------------------------+
| Tenant Status                  | tenant-dr-success                     |
+--------------------------------+---------------------------------------+
| Is Tenant config Refreshed     | false                                 |
+--------------------------------+---------------------------------------+
| Tenant Duration                | 129.163088ms                          |
+--------------------------------+---------------------------------------+
| Device Update Count            | 2                                     |
+--------------------------------+---------------------------------------+
| Device Update Total Duration   | 2m30.478250751s                       |
+--------------------------------+---------------------------------------+
| Maintenance Mode Disable       |                                       |
| Duration                       |                                       |
+--------------------------------+---------------------------------------+
| Start Time                     | 2024-05-15 15:05:02 +0530 IST         |
+--------------------------------+---------------------------------------+
| Last Modified                  | 2024-05-15 15:08:47 +0530 IST         |
+--------------------------------+---------------------------------------+
| Duration                       | 3m45.7253427s                         |
+--------------------------------+---------------------------------------+


Inventory Service Response:
.
.
.

Policy Service Response:
```

```
Tenant service Response:
--- Time Elapsed: 15.230163ms ---
```

SLX: SW-22# show runn router bgp router bgp local-as 4200065535 capability as4-enable fast-external-fallover neighbor 10.20.20.5 remote-as 4200065535 neighbor 10.20.20.5 next-hop-self address-family ipv4 unicast network 172.31.254.211/32 network 172.31.254.248/32 maximum-paths 8 graceful-restart ! address-family ipv6 unicast ! address-family l2vpn evpn graceful-restart **retain route-target all** ! ! SW-22#	The "retain route-target all" configuration doesn't get removed from the border-leaf during DRC.

Example Configuration for Non-Clos Multi Rack Fabric

The following is an example configuration for **non-Clos multi rack fabric**:

```
efa fabric show -name fs
Fabric Name: fs, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-success, Fabric Health: Green
+--------------+------+-----------+------------+------------+--------------+-------------+------------------+----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     |
ROLE    | DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------------+--------------+-------------+------------------+----------------+---------+-------+
| 10.64.208.28 | r1   | S28       | 4200000000 | leaf
| provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
| 10.64.208.29 | r1   | S29       | 4200000000 | leaf
| provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
| 10.64.208.23 | r2   | L23       | 4200065535 | borderleaf
| provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
| 10.64.208.22 | r2   | L22       | 4200065535 | borderleaf
| provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
+--------------+------+-----------+------------+------------+--------------+-------------+------------------+----------------+---------+-------+
```

The following examples show the `retain route-target all` configuration behavior in XCO 3.5.0 and XCO 3.6.0 and later in a **non-Clos multi rack fabric**:

- **XCO 3.5.0 Behavior**

On border leaf or leaf devices

XCO CONFIGURED ROUTER BGP
SLX:
L22# show running-config router bgp address-family l2vpn evpn
router bgp
  address-family l2vpn evpn
  graceful-restart
  retain-route-target all
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
  !
!
L22#
	User created OOB configuration
	User removes the **"retain route-target all"**
	L2# config
	Entering configuration mode terminal
	L22(config)# router bgp
	L22(config-bgp-router)# address-family l2vpn evpn
	L22(config-bgp-evpn)# **no retain route-target all**

```
efa inventory device update -ip 10.64.208.22
efa fabric show

Fabric Name: fs, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-success,
Fabric Health: Red
+--------------+------+-----------+------------+------------+--------------+---------------
+------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     |    ROLE    |
DEVICE STATE |   APP STATE   | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------------+--------------+---------------
+------------------+-----------------+---------+-------+
| 10.64.208.28 | r1   | S28       | 4200000000 | leaf       |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.29 | r1   | S29       | 4200000000 | leaf       |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.23 | r2   | L23       | 4200065535 | borderleaf |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.22 | r2   | L22       | 4200065535 | borderleaf |
provisioned  | cfg refreshed | BGPU              |                 | 2       | 1     |
+--------------+------+-----------+------------+------------+--------------+---------------
+------------------+-----------------+---------+-------+

efa fabric debug device drift --device-ip 10.64.208.28 --name fs --reconcile
No Drift Found hence no reconciliation required
--- Time Elapsed: 440.8283ms ---

"retain route-target all" is not configure back as part of reconcile fabric.

XCO CONFIGURED ROUTER BGP
SLX:
L22(config-bgp-evpn)# do show running-config router bgp address-family l2vpn evpn
router bgp
 address-family l2vpn evpn
  graceful-restart
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
 !
!
L22(config-bgp-evpn)#
```

- **Behavior in XCO 3.6.0 and later**

On border-leaf or leaf devices

XCO CONFIGURED ROUTER BGP	**OOB CREATED:**

SLX:
L22# show running-config router bgp address-family l2vpn evpn
router bgp
  address-family l2vpn evpn
  graceful-restart
  retain route-target all
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
  !
!
L22#
	User removes the **"retain route-target all"**
	L2# config
	Entering configuration mode terminal
	L22(config)# router bgp
	L22(config-bgp-router)# address-family l2vpn evpn
	L22(config-bgp-evpn)# **no retain route-target all**

```
efa inventory device update -ip 10.64.208.22
efa fabric show

Fabric Name: fs, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-success,
Fabric Health: Red
+--------------+------+-----------+------------+------------+--------------+---------------
+------------------+----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     |    ROLE    |
DEVICE STATE |   APP STATE   | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------------+--------------+---------------
+------------------+----------------+---------+-------+
| 10.64.208.28 | r1   | S28       | 4200000000 | leaf       |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.29 | r1   | S29       | 4200000000 | leaf       |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.23 | r2   | L23       | 4200065535 | borderleaf |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.22 | r2   | L22       | 4200065535 | borderleaf |
provisioned  | cfg refreshed | BGPU              |                 | 2       | 1     |
+--------------+------+-----------+------------+------------+--------------+---------------
+------------------+----------------+---------+-------+

drc with reconcile.
efa fabric debug device drift --device-ip 10.64.208.22 --name fs --reconcile
  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  | BgpIPV4Network                  |
| Global | cfg-in-sync  | BgpIPV4NetworkGracefulRestart   |
| Global | cfg-in-sync  | BgpL2EVPNNetworkGracefulRestart |
| Global | cfg-refreshed | BgpL2EVPNRetainRtAll           |
| Global | cfg-in-sync  | BgpIPV4NetworkMaxPath           |
+--------+--------------+---------------------------------+
```

```
+-------------+---------+-------+
| CONFIG TYPE | STATUS  | ERROR |
+-------------+---------+-------+
| routerbgp   | Success |       |
+-------------+---------+-------+
$ efa fabric show --name fs

Fabric Name: fs, Fabric Description: , Fabric Type: non-clos, Fabric
Status: configure-success, Fabric Health: Green
+--------------+------+-----------+------------+------------+--------------+-------------
+------------------+----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     |
ROLE    | DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------------+--------------+-------------
+------------------+----------------+---------+-------+
| 10.64.208.28 | r1   | S28       | 4200000000 | leaf
| provisioned  | cfg in-sync | NA              | NA              | 2       | 1     |
| 10.64.208.29 | r1   | S29       | 4200000000 | leaf
| provisioned  | cfg in-sync | NA              | NA              | 2       | 1     |
| 10.64.208.23 | r2   | L23       | 4200065535 | borderleaf
| provisioned  | cfg in-sync | NA              | NA              | 2       | 1     |
| 10.64.208.22 | r2   | L22       | 4200065535 | borderleaf
| provisioned  | cfg in-sync | NA              | NA              | 2       | 1     |
+--------------+------+-----------+------------+------------+--------------+-------------
+------------------+----------------+---------+-------+
```

XCO CONFIGURED ROUTER BGP SLX: L22(config-bgp-evpn)# do show running-config router bgp address-family l2vpn evpn router bgp address-family l2vpn evpn graceful-restart **retain route-target all** neighbor overlay-ebgp-group encapsulation vxlan neighbor overlay-ebgp-group next-hop-unchanged neighbor overlay-ebgp-group activate ! ! L22(config-bgp-evpn)#	**"retain route-target all"** is configured back as part of reconcile fabric.

Example Configuration for Clos Fabric

The following is an example configuration for a **Clos fabric**:

```
efa fabric show --name stage5
Fabric Name: stage5, Fabric Description: , Fabric Stage: 5, 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 |
+--------------+-----------+-----------+-------+------------+--------------+-------------
+------------------+----------------+---------+-------+
| 10.64.208.29 | podstage5 | S29       | 64769 | superspine | provisioned  | cfg in-
sync | NA               | NA              | NA      | 1     |
| 10.64.208.28 | pod3      | S28       | 64512 | spine      | provisioned  | cfg in-
sync | NA               | NA              | NA      | 1     |
```

```
| 10.64.208.22 | pod3      | L22       | 65000 | leaf       | provisioned  | cfg in-
sync | NA                | NA              | 2       | 1     |
| 10.64.208.23 | pod3      | L23       | 66000 | borderleaf | provisioned  | cfg in-
sync | NA                | NA              | 2       | 1     |
+--------------+-----------+-----------+-------+------------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

The following examples show the `retain route-target all` configuration behavior in XCO 3.5.0 and 3.6.0 in a **Clos fabric**:

- **XCO 3.5.0 Behavior**

  On border leaf devices

XCO CONFIGURED ROUTER BGP
SLX:
L22(config-bgp-evpn)# do show running-config router bgp address-family l2vpn evpn
router bgp
  address-family l2vpn evpn
  graceful-restart
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
  !
!
L22(config-bgp-evpn)#
	User Created OOB Configurations: "retain route-target all"
	L2# config
	Entering configuration mode terminal
	L22(config)# router bgp
	L22(config-bgp-router)# address-family l2vpn evpn
	L22(config-bgp-evpn)# **retain route-target all**

```
efa inventory device update -ip 10.64.208.23
efa fabric show

Fabric Name: stage5, Fabric Description: , Fabric Stage: 5, 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 |
+--------------+-----------+-----------+-------+------------+--------------
+---------------+-------------------+-----------------+---------+-------+
| 10.64.208.29 | podstage5 | S29       | 64769 | superspine |
provisioned  | cfg in-sync   | NA                | NA              | NA      | 1     |
| 10.64.208.28 | pod3      | S28       | 64512 | spine      |
provisioned  | cfg in-sync   | NA                | NA              | NA      | 1     |
| 10.64.208.22 | pod3      | L22       | 65000 | leaf       |
provisioned  | cfg in-sync   | NA                | NA              | 2       | 1     |
| 10.64.208.23 | pod3      | L23       | 66000 | borderleaf |
provisioned  | cfg refreshed | BGPU              |                 | 2       | 1     |
+--------------+-----------+-----------+-------+------------+--------------
+---------------+-------------------+-----------------+---------+-------+

efa fabric configure --name stage5
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: 30.549013784s ---

"retain route-target all" is removed as part of configure fabric.
SLX:
L22# do show running-config router bgp address-family l2vpn evpn
router bgp
 address-family l2vpn evpn
  graceful-restart
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
 !
!
L22(config-bgp-evpn)#
```

- **XCO 3.6.0 Behavior**

  On border-leaf or leaf devices

SLX: L23(config-bgp-evpn)# do show running-config router bgp address-family l2vpn evpn router bgp address-family l2vpn evpn graceful-restart neighbor overlay-ebgp-group encapsulation vxlan neighbor overlay-ebgp-group next-hop-unchanged neighbor overlay-ebgp-group activate ! ! L23(config-bgp-evpn)#	**User Created OOB Configurations: "retain route-target all"** L2# config Entering configuration mode terminal L22(config)# router bgp L22(config-bgp-router)# address-family l2vpn evpn L22(config-bgp-evpn)# **retain route-target all**

```
efa inventory device update -ip 10.64.208.23
efa fabric show

Fabric Name: stage5, Fabric Description: , Fabric Stage: 5, 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 |
+--------------+-----------+-----------+-------+------------+--------------
+---------------+------------------+-----------------+---------+-------+
| 10.64.208.29 | podstage5 | S29       | 64769 | superspine |
provisioned  | cfg in-sync   | NA              | NA              | NA      | 1     |
| 10.64.208.28 | pod3      | S28       | 64512 | spine      |
provisioned  | cfg in-sync   | NA              | NA              | NA      | 1     |
| 10.64.208.22 | pod3      | L22       | 65000 | leaf       |
provisioned  | cfg in-sync   | NA              | NA              | 2       | 1     |
| 10.64.208.23 | pod3      | L23       | 66000 | borderleaf |
provisioned  | cfg refreshed | BGPU            |                 | 2       | 1     |
+--------------+-----------+-----------+-------+------------+--------------
+---------------+------------------+-----------------+---------+-------+



efa fabric configure --name stage5
```

```
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: 30.549013784s ---


"retain route-target all" is not removed as part of configure fabric that is expected.
XCO CONFIGURED ROUTER BGP
SLX:
L23(config-bgp-evpn)# do show running-config router bgp address-family l2vpn evpn
router bgp
 address-family l2vpn evpn
  graceful-restart
  retain route-target all
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
 !
!
L22(config-bgp-evpn)#
efa fabric debug device drift --device-ip 10.64.208.28 --name stage5 --reconcile
No Drift Found hence no reconciliation required
--- Time Elapsed: 440.8283ms ---

"retain route-target all" is not removed as part of reconcile, that is expected.
XCO CONFIGURED ROUTER BGP
SLX:
L23(config-bgp-evpn)# do show running-config router bgp address-family l2vpn evpn
router bgp
 address-family l2vpn evpn
  graceful-restart
  retain route-target all
  neighbor overlay-ebgp-group encapsulation vxlan
  neighbor overlay-ebgp-group next-hop-unchanged
  neighbor overlay-ebgp-group activate
 !
!
L22(config-bgp-evpn)#
```

Configure SLX Password Expiry Notification

You can configure SLX password expiry notification.

About This Task

Follow this procedure to configure SLX password expiry notification.

- XCO 3.5.0 and later enables you to configure SLX-OS user password attributes for password expiry, and is supported from SLX version 20.6.1.
- There is no default configuration for the SLX-OS user password attributes on SLX.
- For each user, alert and alarm will be generated when their password expires.
- The password expiry settings apply to all users except the root user.

- The retiring user will not generate any alerts and will not clear the health.
- User management will be directly done on SLX devices.

Procedure

Run the following command to configure SLX password expiry notification:

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

```
# efa inventory device setting update --ip 10.64.192.72 --password-expiry-info 30 --password-expiry-
minor 20 --password-expiry-major 10 --password-expiry-critical 2
+--------------+--------------------------------+---------+-------+-------+
|  IP ADDRESS  |              NAME              | STATUS  | VALUE | ERROR |
+--------------+--------------------------------+---------+-------+-------+
| 10.64.192.72 | Password Expiry For            | Success | 30    |       |
|              | Info Level                     |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for            | Success | 20    |       |
|              | Minor Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for            | Success | 10    |       |
|              | Major Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for            | Success | 2     |       |
|              | Critical Level                 |         |       |       |
+--------------+--------------------------------+---------+-------+-------+

--- Time Elapsed: 53.849675ms ---

# efa inventory device setting update --fabric myclos --password-expiry-info 30 --password-expiry-
minor 20 --password-expiry-major 10 --password-expiry-critical 2
+--------------+--------------------------------+---------+-------+-------+
|  IP ADDRESS  |              NAME              | STATUS  | VALUE | ERROR |
+--------------+--------------------------------+---------+-------+-------+
| 10.64.192.72 | Password Expiry For            | Success | 30    |       |
|              | Info Level                     |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for            | Success | 20    |       |
|              | Minor Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for            | Success | 10    |       |
|              | Major Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for            | Success | 2     |       |
|              | Critical Level                 |         |       |       |
+--------------+--------------------------------+---------+-------+-------+

--- Time Elapsed: 53.849675ms ---


# efa inventory device setting show --ip 10.64.192.72
+--------------------------------+-------+-------------+
|              NAME              | VALUE |  APP STATE  |
+--------------------------------+-------+-------------+
| Password Expiry Info           | 90    |             |
+--------------------------------+-------+-------------+
| Password Expiry Minor          | 60    |             |
+--------------------------------+-------+-------------+
| Password Expiry Major          | 30    |             |
+--------------------------------+-------+-------------+
| Password Expiry Critical       | 2     |             |
+--------------------------------+-------+-------------+
```

```
|  Max-password-age              | 90    |             |
+--------------------------------+-------+-------------+
```

SLX Password Expiry Alerts

Password expiry alerts are generated in advance of the expiration date, based on certain predefined days. There are different levels of alerts, such as info level days, minor level days, and major level days.

The following are the SLX password expiry alerts:

- **PasswordExpiryThresholdAlert**: It is raised when the password expiry date has reached a certain threshold, such as 90 days, 60 days, or 30 days. These expiry days can be set using a configuration command.
- **PasswordExpiredAlert**: It is generated on the day when a user's password has expired, and is generated every day until the password is renewed.
- **PasswordExpiryClearAlert**: It is raised when the user has successfully renewed their password.

For more details on SLX password expiry alerts, see the "Password Expiry Alerts Inventory" table in the topic Inventory of Alerts on page 925.

BFD Session Transition on IP Fabric Underlay

Learn about architecture and step-by-step provisioning guidance for executing the BFD session migration safely, without disrupting BGP adjacencies or causing routing instability.

Overview

ExtremeCloud Orchestrator (XCO) 4.1.0 introduces the capability to transition BFD (Bidirectional Forwarding Detection) sessions on IP fabric underlay Layer 3 interfaces between hardware-accelerated and software-terminated operation modes, with full support for the reverse direction. This capability enables safe migration of routing-critical BFD sessions while maintaining routing stability and minimizing traffic impact.

BFD provides fast, sub-second liveliness detection for routing adjacencies. XCO 4.1.0 lets you choose how each routing-relevant BFD session is terminated:

- **Hardware-based BFD** — BFD packets are processed in the device data plane.
- **Software-based BFD** — BFD packets are processed by the control plane so that session state reflects the true state of the routing control plane.

The transition migrates all routing-relevant BFD sessions — Leaf-to-Spine, Spine-to-Super-Spine, and Leaf-to-Leaf over the ICL/MCT — from one termination type to the other. The migration is topology-aware and is designed to complete without resetting eBGP or iBGP adjacencies.

BFD (Bidirectional Forwarding Detection) provides rapid failure detection for routing protocols such as BGP. XCO supports controlled migration in both directions:

- Hardware-based BFD to software-based BFD

- Software-based BFD to hardware-based BFD

This migration is performed using a topology-aware workflow that prevents widespread service disruption and preserves BGP adjacencies during the transition.

Feature Highlights

- **Safe, non-disruptive migration** — Migrates all routing-critical BFD sessions (underlay eBGP, iBGP, and ICL/MCT) between hardware-based and software-based termination without triggering BGP session resets or routing-table churn.
- **New user-configurable BFD session type setting** — Introduces the --bfd-session-type {hardware|software} fabric-level setting, giving operators explicit, auditable control over the BFD termination plane.
- **Green-field fabric support** — Allows the desired BFD session type to be declared at fabric creation time, ensuring new deployments are provisioned correctly from the first configure operation.
- **In-service transition on active fabrics** — Supports on-the-fly BFD session-type changes on live fabrics without impacting existing BGP adjacencies or causing interface flaps.
- **Topology-aware migration engine** — Implements distinct migration strategies for Clos (spine-by-spine, plane-scoped) and non-Clos (rack-by-rack) topologies, guaranteeing a deterministic and bounded blast radius at every step.

Deliverables and Requirements

Requirement	Description
Software-Terminated BFD	All BFD sessions that influence routing decisions—eBGP peer-group sessions on underlay Ethernet interfaces and iBGP sessions over ICL/MCT port-channels—MUST be terminated in software upon completion of migration.
Zero Routing Adjacency Disruption	The migration workflow MUST NOT cause eBGP or iBGP session resets, route withdrawals, or BGP hold-timer expirations at any point during execution.
Deterministic Blast Radius	At any instant during migration, BFD policy is disabled on at most one well-defined topology unit: exactly one Spine plane (CLOS) or exactly one Rack/Logical Leaf (Non-CLOS). Concurrent multi-unit migration is prohibited
Complete Adjacency Coverage	Migration scope MUST encompass all routing-relevant BFD sessions: Leaf « Spine, Spine « SuperSpine, and Leaf « Leaf (ICL/MCT iBGP).
Automation-Readiness	The end-to-end workflow MUST be deterministic, topology-aware, and fully enforceable through the XCO automation engine without manual per-device intervention.

Limitations

Platform Exclusion: BFD session-type transition is not supported on OS ONE-platform devices. Unsupported devices are excluded from software BFD migration processing.

Note

- Fabric service doesn't support rollback. You must either trigger 'fabric configure' to continue with the current update or revert settings to the previous state and then do 'fabric configure'
- Upgrading XCO from 4.0.2 to 4.1.0 does not change BFD behavior. All existing BFD sessions remain hardware-based, no automatic transition is triggered, and the fabric continues to operate with hardware BFD. Transitioning to software BFD is an explicit, operator-controlled action — XCO never performs an implicit transition.

Platforms Supported

Requires firmware baseline: SLX-OS 20.8.2.

Note
Upgrade the device firmware before enabling software BFD on a fabric.

Platform	Software BFD Support	Notes
Extreme 8520	Supported	Full hardware and software BFD
SLX 9540 / 9640	Not Supported	Hardware BFD offload not available on this platform for fabric underlay
SLX 9740	Supported	Full hardware and software BFD
Extreme 8820	Supported	Full hardware and software BFD
Extreme 8720	Supported	Full hardware and software BFD
Extreme OS ONE Platforms	Not Supported	NA

Why Hardware-Based BFD Is Unsafe in Routing-Critical Paths

In a network operating hardware-offloaded BFD, liveness detection occurs at the ASIC (Application-Specific Integrated Circuit) level, entirely below the operating system's routing stack. This architecture delivers sub-50 ms failure detection but decouples liveness signaling from control-plane health. The unsafe failure modes are:

- **Control-plane stall:** Routing daemons (BGP, route manager) may pause due to garbage collection, high CPU load, or software faults. The ASIC continues answering BFD Hellos, so the remote peer sees no failure, and BGP holds the adjacency open.

- **Inability to withdraw routes:** If the routing daemon cannot process route-withdraw requests, traffic continues to be forwarded to the stalled node. The node cannot drain the traffic, producing a blackhole.
- **Silent failure propagation:** Because no BFD-down event is generated, upstream nodes and load-balancers receive no signal to reroute traffic, extending the outage duration beyond what hold-timer expiry would produce.

Software-terminated BFD restores control-plane coupling: when the routing daemon pauses, BFD Hello transmissions stop, the hold timer expires on the remote peer, the BFD session transitions to DOWN, and BGP withdraws the associated routes within milliseconds.

Benefits of Software-Based BFD

Hardware-based BFD can remain operational even when the routing control plane becomes unresponsive. In such situations, traffic may continue to be forwarded toward devices that can no longer properly participate in routing convergence, potentially resulting in traffic blackholing.

Software-based BFD addresses this limitation by associating BFD liveliness with the control-plane state, providing:

- Improved detection of routing process failures
- Better alignment between failure detection and routing convergence
- Reduced risk of silent traffic blackholing
- Controlled and deterministic migration workflows

Note
The guiding principle is simple: a node must not keep forwarding traffic unless it can also withdraw routes. Software-based BFD makes the session reflect control-plane state, which restores this behavior.

Policy–Mechanism Separation

The migration design formally separates two orthogonal dimensions of BFD configuration:

Dimension	Configuration Object	What It Controls
BFD Policy	`Ethernet or port-channel interface`	Whether BFD is administratively enabled on a specific underlay interface. Disabling the policy brings the BFD session to admin-down without signaling a protocol failure to BGP.
BFD Mechanism	`bfd-software-session (interface keyword) / --bfd-session-type (fabric setting)`	Whether BFD is processed by hardware (ASIC) or software (control-plane daemon). Changing the mechanism requires the session to be re-negotiated.

For every migration unit, XCO applies the following sequence, which guarantees zero BGP adjacency churn:

1. Administratively disable BFD policy on all underlay Ethernet interfaces and port-channels belonging to the current migration unit (Spine plane or Rack).
2. Modify the BFD mechanism on those interfaces from hardware to software (or vice versa).Re-enable BFD policy on all affected interfaces.
3. Re-enable BFD policy on all affected interfaces or ports.

Note
When BFD is disabled on a peer group, the session goes down with reason admin_down. Because the peers signal an administrative down (not a protocol failure), BGP is not brought down, and adjacencies stay UP throughout the transition.

Interface-Level vs BGP-Level BFD Policy

XCO applies BFD policy at the underlay interface level rather than attaching it directly to the BGP neighbor statement. This design choice has the following rationale:

- **Protocol independence:** Interface-level BFD detects data-plane link failures regardless of which control-plane protocols are running over the interface (BGP, OSPF, IS-IS, static routes). A single BFD session protects all protocols simultaneously.
- **Data-plane visibility:** Interface BFD detects forwarding-path failures that do not manifest as routing-protocol events (e.g., unidirectional fiber failures, ASIC forwarding-table inconsistencies).
- **Consistent adjacency semantics:** All routing protocols using the interface react to the BFD-down event simultaneously, preventing split-brain scenarios where BGP withdraws routes but OSPF or IS-IS continues advertising reachability.

BFD Application Point	Failure Scope	Protocols Notified
BGP neighbor statement	BGP session only	BGP only
Underlay Ethernet / port-channel interface (XCO model)	Physical link and forwarding path	All protocols bound to the interface (BGP, OSPF, IS-IS, statics)

BFD Migration Strategies (Horizontal vs Vertical Migration)

When migrating BFD from a hardware-based implementation to a software-based implementation, two migration strategies can be used:

- **Horizontal BFD Migration**
- **Vertical BFD Migration**

Horizontal BFD Migration

Horizontal migration introduces a temporary state in which BFD protection is unavailable across the entire fabric. This can delay failure detection, slow network convergence, and increase the risk of traffic loss during the migration window.

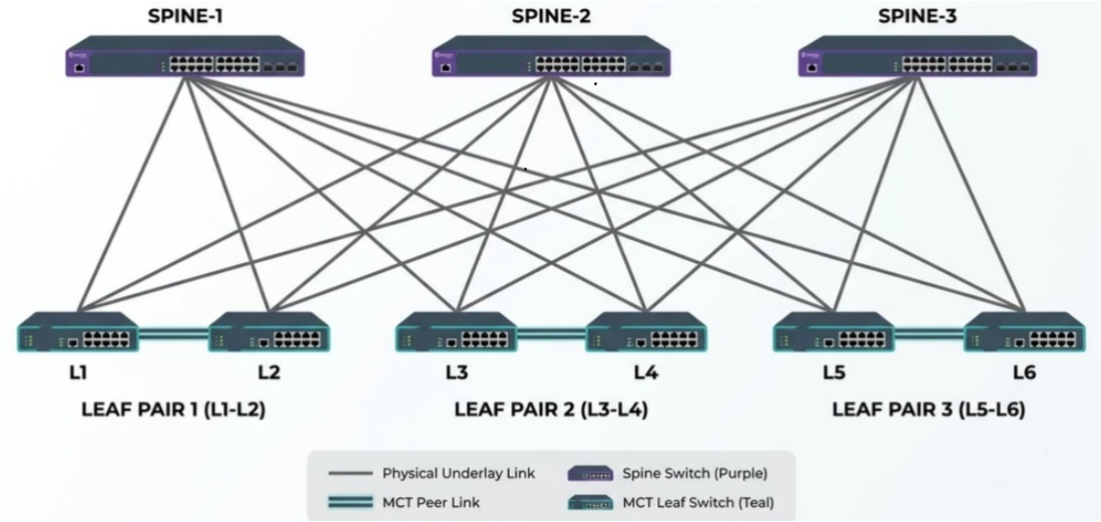

Migration Workflow

Perform the following tasks:

- Disable the BFD policy on all underlay interfaces and underlay port-channels.
- Transition BFD from the hardware implementation to the software implementation
- Re-enable the BFD policy on all interfaces.

Operational Behavior

During a horizontal migration:

- BFD is administratively disabled on all underlay links and port-channels across the fabric.
- Software-based BFD is enabled on all participating interfaces.
- After the migration is complete, BFD is re-enabled throughout the fabric.

Failure Handling

The migration is executed on a device-by-device basis.

For example:

- Spine-1
- Spine-2
- Spine-3

If a migration failure occurs on any device:

- The migration process stops immediately.
- An error message is generated and reported to the user.
- Manual intervention may be required before restarting the migration.

Network Impact

Because BFD is disabled fabric-wide during the migration:

- BGP loses fast failure detection capabilities.
- BGP relies on keepalive and hold-timer mechanisms to detect failures.
- Link or path failures may remain undetected until hold timers expire.

As a result:

- Failure detection time increases from milliseconds to several seconds.
- Network convergence is significantly delayed.
- Traffic can be forwarded toward failed paths during the detection interval

Scalability Considerations

The operational risk increases as the fabric grows.

In larger deployments:

- More BGP sessions are affected simultaneously.
- The probability of undetected failures increases during migration.
- Traffic impact becomes more pronounced.
- The potential for temporary traffic blackholing increases.

Note
During the period in which BFD is disabled across the fabric, failure detection relies entirely on BGP timers. Any forwarding failure occurring during this interval may not be detected immediately

Vertical BFD Migration

In a vertical migration, BFD is migrated in stages. Only a specific fabric segment or plane is migrated at a time, while the remainder of the fabric continues to operate with BFD protection enabled.

Vertical migration preserves BFD protection throughout most of the fabric by migrating one segment at a time. This staged approach minimizes traffic impact, maintains fast failure detection, and supports predictable network convergence during migration.

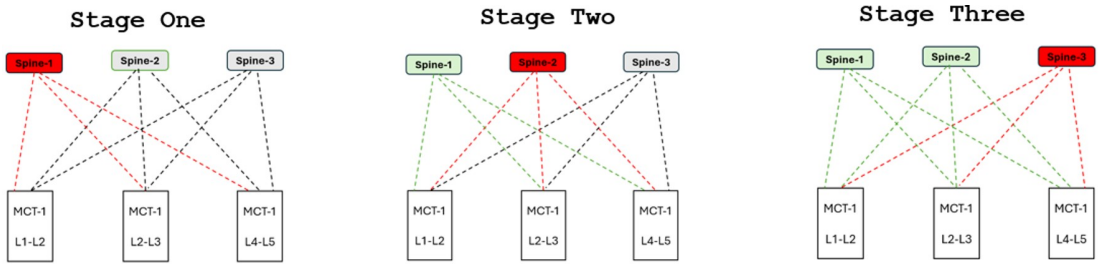

Migration Workflow

Perform the following tasks for each migration segment:

- Disable the BFD policy for the target plane or fabric segment.
- Transition BFD from the hardware implementation to the software implementation.
- Re-enable the BFD policy for the migrated segment.

Repeat the process for the next segment.

Migration Strategy

Vertical migration uses a phased approach that minimizes the scope of change during each stage.

- Clos Fabrics (Three-Stage)
  - Migrate one spine plane at a time.
  - Validate stability before proceeding to the next spine plane.
- Five-Stage Fabrics
  - Migrate spine and superspine planes sequentially.
  - Verify successful operation after each migration stage.
- Non-Clos Fabrics
  - Migrate the fabric on a rack-by-rack basis.
  - Limit migration impact to a single rack segment at a time.
- MCT Leaf Pairs
  - Complete spine-plane stabilization first.
  - Migrate MCT leaf pairs in parallel after validation.

Failure Handling

If a migration failure occurs during any stage:

- The migration process stops immediately.
- The remaining fabric continues to operate with BFD enabled.
- Error information is reported to the user.
- The issue remains isolated to the affected migration segment.

Continuous Failure Protection

A key advantage of vertical migration is that BFD protection remains active across most of the fabric throughout the migration.

At any given time:

- Only a small portion of the fabric has BFD temporarily disabled.
- The majority of network paths continue to benefit from fast failure detection.
- Most traffic flows remain protected by existing BFD sessions.

Network Impact

Even when a specific plane or segment is being migrated:

- Other spine planes continue detecting failures within milliseconds.
- BGP convergence remains fast across unaffected network paths.
- Failure isolation is maintained.

This approach helps prevent:

- Traffic blackholing
- Delayed failure detection
- Fabric-wide instability
- Large-scale convergence delays

Benefits of Vertical Migration

Vertical migration provides the following advantages:

- Eliminates the need for a fabric-wide BFD outage.
- Maintains fast failure detection across most traffic paths.
- Limits operational risk to a small migration domain.
- Reduces the likelihood of traffic disruption.
- Supports controlled execution and safer rollback procedures.
- Improves overall migration reliability in large-scale fabrics.

Figure 15:

Side-by-Side Comparison

Attribute	Horizontal migration	Vertical (staged) migration
Migration scope	Entire fabric in one pass	One plane / segment at a time
BFD disabled window	**Fabric-wide**	**Single segment only**
Failure detection during migration	Falls back to BGP hold timer (seconds)	Milliseconds on all non-migrating segments
Blackhole / traffic-loss risk	**High**	**Low (bounded)**
Impact as fabric scales	Worsens with size	Stays bounded
Rollback control	Coarse (whole fabric)	Fine (per segment)
On failure	Aborts; fabric was globally exposed	Stops; rest of fabric stays protected

Tip
Vertical (staged) BFD migration is the preferred approach because it avoids global BFD disablement, ensures continuous failure detection, and prevents traffic loss during migration. Horizontal migration should be avoided on production fabrics because it exposes the entire fabric to delayed failure detection for the duration of the change.

Fabric Setting Update — BFD Session Type Flag

The BFD session type is a fabric-level setting configured through the efa fabric setting update command. The relevant flag is:

```
--bfd-session-type <STRING: hardware | software> (default: hardware)
```

Setting this flag to `software` instructs XCO to provision the `bfd-software-session` keyword on all routing-relevant underlay interfaces during the next **`efa fabric configure`** command execution. The flag also drives the staged migration workflow when applied to an already-provisioned active fabric.

The following table describes the key CLI flags relevant to BFD configuration:

Flag	Type / Range	Description
`--bfd-enable`	`Yes | No`	Enables or disables BFD globally for the fabric underlay.
`--bfd-tx`	`50–30000 ms (SLX) 10–30000 ms (Extreme OS)`	Desired minimum BFD transmit interval.
`--bfd-rx`	`50–30000 ms (SLX) 10–30000 ms (Extreme OS)`	Desired minimum BFD receive interval.
`--bfd-multiplier`	`3–50`	BFD detection-time multiplier. Detection time = multiplier × max(tx, rx).
`--bfd-session-type`	`hardware | software`	Controls ASIC-offloaded vs control-plane BFD termination. Default: hardware.

For more information on the fabric settings flags, see the *Fabric Settings in Active Fabric: Small Data Center Fabric* and *Fabric Settings in Active Fabric: Clos Fabric* tables in the Edit Fabric Settings on page 212 topic.

Device Running-Configuration Output

On a Spine device provisioned with software BFD, the interface-level running configuration includes the `bfd-software-session` keyword:

```
SLX-30# show running-config interface Ethernet 0/3-5
interface Ethernet 0/3
 bfd-software-session
 ip address 10.10.10.19/31
 bfd interval 300 min-rx 300 multiplier 3
 no shutdown
!
interface Ethernet 0/4
 bfd-software-session
 ip address 10.10.10.23/31
 bfd interval 300 min-rx 300 multiplier 3
 no shutdown
!
interface Ethernet 0/5
 bfd-software-session
 ip address 10.10.10.21/31
 bfd interval 300 min-rx 300 multiplier 3
 no shutdown
!
```

The absence of `bfd-software-session` on an interface indicates that hardware BFD is active (default behaviour). The BGP peer-group configuration is unchanged—BFD attachment at the peer-group level only determines whether BFD is consulted; the termination plane is governed solely by the interface keyword.

XCO Automation, Versioning, and Upgrade

XCO ensures the following:

- **Non-disruptive upgrades**

  BFD behavior remains unchanged unless explicitly modified.
- **Operator-controlled transition**

  Migration to Software BFD is fully opt-in.
- **Controlled migration model**

  All transitions follow:
  - Strict safety rules
  - Bounded blast-radius principles
  - Topology-aware staged execution
- **Bidirectional Transition Supported**

  Users can revert from software-based BFD back to hardware-based BFD.
- **No Automatic Mode**

  There is no auto-selection mechanism; hardware BFD remains the default.
- **Scope of Impact**
  - Applies only to **underlay BFD sessions**
  - **Overlay (CEP) connections are not affected**

This structured approach ensures predictable, safe, and controlled migration of BFD session types across supported fabric deployments.

XCO Automation Model

ExtremeCloud Orchestrator (XCO) enforces a deterministic and safeguards-driven automation model to ensure safe and consistent fabric operations during migration workflows.

XCO performs the following actions:

- **Topology Detection:** Automatically identifies the fabric topology (Clos or Non-Clos and adapts the workflow accordingly.
- **Migration Unit Enforcement:** Applies the correct migration granularity
  - **Clos fabrics** — Spine-based migration
  - **Non-Clos fabrics** — Rack-based migration

- **Serialized Execution:** Prevents concurrent or parallel migrations to maintain operational stability.
- **Super-Spine Workflow Restriction:** SuperSpine-specific migration workflows are blocked unless SuperSpine is active via BorderLeaf connectivity, keeping migrations in line with supported fabric designs.

Note
These controls ensure a bounded blast radius and predictable migration behavior across all supported topologies.

Behavior During Upgrade from XCO 4.0.2 to XCO 4.1.0

When upgrading an existing fabric from XCO 4.0.2 to XCO 4.1.0:

- All existing BFD sessions **remain hardware-based**.
- No automatic transition to Software BFD is triggered.
- The fabric continues to operate using hardware BFD.

This behavior ensures:

- No implicit control-plane changes during upgrade
- Preservation of backward compatibility
- Continued operational stability

Note

- XCO upgrades do not modify BFD session behavior unless explicitly configured by the operator.
- The transition to software BFD is opt-in and operator-controlled. XCO never performs an implicit or automatic BFD session transition. You must explicitly change the fabric setting and run **`efa fabric configure`** to start the staged, topology-aware migration.
- **Plane-scoped Software BFD migration** is supported starting from XCO version 4.1.0 and later.

Migration Behavior

After configuration is applied, XCO initiates a **staged, topology-aware migration workflow**:

- **Clos fabrics**
  - Spine-by-spine migration
  - Plane-scoped transition
- **Non-Clos fabrics**
  - Rack-by-rack migration

Note
XCO does not perform any automatic or implicit BFD transition without this configuration change.

Operational Impact

When the fabric setting is updated:

- Fabric state transitions to **`settings-updated`**
- Fabric health may temporarily appear **non-green**

During migration:

- BFD sessions may flap
- **BGP sessions remain stable**, ensuring no underlay disruption

Monitoring and Failure Handling

- Migration progress aligns with staged configuration workflows (similar to MD5 updates)
- If configuration fails:
  - `efa fabric configure` reports failure
- BFD-specific issues are:
  - Reported as **warnings**, not fatal errors

Note
No additional migration-specific states are exposed in `efa fabric show`.

User-Initiated Transition to Software BFD

After upgrading to XCO 4.1.0, transitioning from hardware-based BFD to software-based BFD requires explicit user action.

Before You Begin

- Fabric must be upgraded to **XCO 4.1.0 or later**
- Supported platforms include:
  - SLX 9740, Extreme 8820, Extreme 8720
  - Extreme 8720, Extreme 8520 (platform-dependent support)

Note
OS ONE currently does not support BFD session type configuration.

About This Task
Follow this procedure to transition from hardware-based BFD to software-based BFD.

Procedure

1. Configure the fabric to use software-based BFD.

   ```
   bfd-session-type software
   ```

   Default value: hardware
2. Apply the configuration.

   ```
   efa fabric configure
   ```

Provision Fabric with BFD Configuration

You can provision fabric with BFD configuration.

About This Task

Follow this procedure to provision the fabric with BFD configuration.

The following end-to-end examples in this procedure use a 3-stage Clos fabric named fs with one Spine (SLX-30 @ 10.64.160.30) and two Leafs (SLX-11 @ 10.64.160.11, SLX-38 @ 10.64.160.38).

Note

- The default BFD session type is hardware. A transition always requires both an **explicit setting update** and an **explicit efa fabric configure**.
- If a device runs firmware that does not support software BFD, it is correctly skipped during staged migration. If that device is later upgraded to a supporting firmware, the software-BFD recipe is **not** automatically re-pushed during device re-configuration, so the device can remain on hardware/default BFD and out of sync with the fabric setting. After upgrading such a device, re-run the transition to bring it into sync.

Procedure

1. Configure software-based BFD on a new fabric.

   You can create a new (green-field) fabric and configure it with software-terminated BFD from inception.

   a. Create the fabric.

   Initializes the fabric record in XCO with default settings. Default BFD session type: hardware.

```
efa fabric create --name fs
Create Fabric fs1 [Success]
--- Time Elapsed: 108.514623ms ---
```

   b. Override BFD session type before adding devices.

   Records the software-BFD intent. No device changes yet.

```
efa fabric setting update --name "fs" -bfd-session-type software
fs Fabric Update Successful
--- Time Elapsed: 4.41886692s ---
```

   c. Add fabric devices.

   Registers and inventories devices. Fabric health transitions to Red (not provisioned).

```
efa fabric device add-bulk --name fs --spine 10.64.160.30 --leaf
10.64.160.11,10.64.160.38  --username admin --password <password>
Inventory Device(s) Registration[Success]
+----+--------------+-----------+-------+--------------+--------------+---------
+--------+
| ID |  IP Address  | Host Name | Model | Chassis Name |   Firmware
| Status  | Reason |
+----+--------------+-----------+-------+--------------+--------------+---------
+--------+
| 23 | 10.64.160.11 | SLX-11    | 3012  | Extreme 8720-32C |
20.8.1_swBfd | Success |    |
+----+--------------+-----------+-------+--------------+--------------+---------
+--------+
```

```
| 27 | 10.64.160.38 | SLX-38    | 3012  | Extreme 8720-32C |
20.8.1_swBfd | Success |    |
+----+--------------+-----------+-------+--------------+--------------+---------
+--------+
| 25 | 10.64.160.30 | SLX-30    | 3012  | Extreme 8720-32C |
20.8.1_swBfd | Success |    |
+----+--------------+-----------+-------+--------------+--------------+---------
+--------+
Device Details
Add Device(s) [Success]

        Addition of Spine device with ip-address = 10.64.160.30 [Succeeded]
        Addition of Leaf device with ip-address = 10.64.160.38 [Succeeded]
        Addition of Leaf device with ip-address = 10.64.160.11 [Succeeded]
Validate Fabric [Success]
--- Time Elapsed: 2m49.489256049s ---
```

d. Configure the fabric.

XCO provisions all devices. Because bfd-session-type is software, the bfd-software-session keyword is applied to all underlay interfaces. Expected result: configure-success, Health: Green.

```
efa fabric configure  --name fs
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: 48.527742952s ---

efa fabric show -name fs
Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+-------------+-------------
+------------------+-----------------+---------+-------+
| 10.64.160.30 |     | SLX-30    | 64512 | Spine | provisioned
| cfg in-sync | NA                | NA             | NA      | 1     |
| 10.64.160.11 |     | SLX-11    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                | NA             | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                | NA             | 2       | 1     |
+--------------+-----+-----------+-------+-------+-------------+-------------
+------------------+-----------------+---------+-------+
```

e. Verify device configuration and BFD sessions.

Check the interface configuration. On a device, it confirms that bfd-software-session is present on the Spine's underlay interfaces.

```
# show running-config interface Ethernet 0/3-5
interface Ethernet 0/3
 bfd-software-session
 ip address 10.10.10.19/31
 bfd interval 300 min-rx 300 multiplier 3
 no shutdown
!
interface Ethernet 0/4
 bfd-software-session
 ip address 10.10.10.23/31
 bfd interval 300 min-rx 300 multiplier 3
 no shutdown
!
interface Ethernet 0/5
```

```
 bfd-software-session
 ip address 10.10.10.21/31
 bfd interval 300 min-rx 300 multiplier 3
 no shutdown
!
```

The presence of bfd-software-session indicates software-based BFD operation.

f. Confirm the sessions are UP and terminated in software. In show bfd neighbors detail, the Session Type field reads Software:

```
# show bfd neighbors detail
Flags: * indicates State is inconsistent across the cluster
OurAddr                                  NeighAddr
State     Int             Session-type
=======                                  =========
=====     ===             ============
10.10.10.18                               10.10.10.19
UP       Eth 0/2         SingleHop

  Local    State: UP            Diag: 0              Demand mode: 0    Poll: 0
  Received State: UP            Diag: 0              Demand mode: 0    Poll: 0
Final: 0
  Local    MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Received MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Rx Count: 5                   Tx Count: 6
  LD/RD:      16012/11          Heard from Remote: Y
  Current Registered Protocols: bgp
  Uptime: 0 day 2 hour 55 min 53 sec 776 msec
  Session Type: Software

OurAddr                                  NeighAddr
State     Int             Session-type
=======                                  =========
=====     ===             ============
10.10.10.20                               10.10.10.21
UP       Eth 0/5         SingleHop

  Local    State: UP            Diag: 0              Demand mode: 0    Poll: 0
  Received State: UP            Diag: 0              Demand mode: 0    Poll: 0
Final: 0
  Local    MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Received MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Rx Count: 5                   Tx Count: 6
  LD/RD:      16012/11          Heard from Remote: Y
  Current Registered Protocols: bgp
  Uptime: 0 day 2 hour 55 min 53 sec 776 msec
  Session Type: Software

OurAddr                                  NeighAddr
State     Int             Session-type
=======                                  =========
=====     ===             ============
10.10.10.22                               10.10.10.23
UP      Eth 0/4         SingleHop

  Local    State: UP            Diag: 0              Demand mode: 0    Poll: 0
  Received State: UP            Diag: 0              Demand mode: 0    Poll: 0
Final: 0
  Local    MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Received MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Rx Count: 5                   Tx Count: 6
  LD/RD:      16012/11          Heard from Remote: Y
  Current Registered Protocols: bgp
```

```
  Uptime: 0 day 2 hour 55 min 53 sec 776 msec
  Session Type: Software
```

g. Use the following RASlog message IDs to confirm session state changes on the SLX devices:

Message ID	Session state
BFD-1001	**UP**
BFD-1002	**DOWN**

The following is a sample log traces:

```
[BFD-1001], 409272,, INFO, SLX,  BFD Session UP for neighbor 10.20.20.3 on
interface Port-channel 64
[BFD-1002], 409214,, INFO, SLX,  BFD Session DOWN for neighbor 10.20.20.3 on
interface Port-channel 64 reason Path Down
```

2. Verify BFD operational state.

To view the BFD neighbors, run the **`show bfd neighbors`** command.

```
SLX# show bfd neighbors
Flags: * indicates State is inconsistent across the cluster
OurAddr                                 NeighAddr
State     Int             Session-type
=======                                 =========
=====     ===             ============
10.10.10.18                             10.10.10.19
UP        Eth 0/2         SingleHop
10.10.10.20                             10.10.10.21
UP        Eth 0/5         SingleHop
10.10.10.22                             10.10.10.23
UP        Eth 0/4         SingleHop

SLX# show bfd neighbors detail
Flags: * indicates State is inconsistent across the cluster
OurAddr                                 NeighAddr
State     Int             Session-type
=======                                 =========
=====     ===             ============
10.10.10.18                              10.10.10.19
UP       Eth 0/2         SingleHop

  Local    State: UP            Diag: 0              Demand mode: 0    Poll: 0
  Received State: UP            Diag: 0              Demand mode: 0    Poll: 0
Final: 0
  Local    MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Received MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Rx Count: 5                   Tx Count: 6
  LD/RD:      16012/11          Heard from Remote: Y
  Current Registered Protocols: bgp
  Uptime: 0 day 2 hour 55 min 53 sec 776 msec
  Session Type: Software

OurAddr                                 NeighAddr
State     Int             Session-type
=======                                 =========
=====     ===             ============
10.10.10.20                              10.10.10.21
UP       Eth 0/5         SingleHop

  Local    State: UP            Diag: 0              Demand mode: 0    Poll: 0
  Received State: UP            Diag: 0              Demand mode: 0    Poll: 0
Final: 0
```

```
  Local    MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Received MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Rx Count: 5                   Tx Count: 6
  LD/RD:      16012/11          Heard from Remote: Y
  Current Registered Protocols: bgp
  Uptime: 0 day 2 hour 55 min 53 sec 776 msec
  Session Type: Software

OurAddr                                 NeighAddr
State     Int             Session-type
=======                                 =========
=====     ===             ============
10.10.10.22                              10.10.10.23
UP      Eth 0/4         SingleHop

  Local    State: UP            Diag: 0              Demand mode: 0    Poll: 0
  Received State: UP            Diag: 0              Demand mode: 0    Poll: 0
Final: 0
  Local    MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Received MinTxInt(ms): 300    MinRxInt(ms): 300    Multiplier: 3
  Rx Count: 5                   Tx Count: 6
  LD/RD:      16012/11          Heard from Remote: Y
  Current Registered Protocols: bgp
  Uptime: 0 day 2 hour 55 min 53 sec 776 msec
  Session Type: Software
```

3. Add devices to software BFD fabric.

   Use this step when the fabric already has software BFD configured and you are adding a new leaf device.

   a. Confirm current fabric state.

      Verify the fabric is already set to software before adding the device.

```
efa fabric setting show --name fs --advanced | grep BFD | BFD Session Type |
software |

efa fabric setting show --name fs -advanced | grep "BFD"
+--------------------------------+-------------------------+
| BFD Enable                     | Yes                     |
+--------------------------------+-------------------------+
| BFD Tx                         | 300                     |
+--------------------------------+-------------------------+
| BFD Rx                         | 300                     |
+--------------------------------+-------------------------+
| BFD Multiplier                 | 3                       |
+--------------------------------+-------------------------+
| BFD Session Type               | hardware                |
+--------------------------------+-------------------------+

--- Time Elapsed: 115.185478ms ---
```

   b. Add the new leaf device.

      Registers the new device. Existing devices transition to cfg-refreshed state. Fabric health changes to Black.

```
efa fabric device add --name "fs" --ip 10.64.160.38 --role leaf --hostname
SLX-160-38
Inventory Device(s) Registration[success]
+----+--------------+-----------+-------+--------------+------------------+
| ID |  IP Address  | Host Name | Model | Chassis Name |   Firmware       |
+----+--------------+-----------+-------+--------------+------------------+
| 5  | 10.64.160.38 | SLX-38    | 3012  | Extreme 8720-32C| 20.8.1_po_bfd |
+----+--------------+-----------+-------+--------------+---------------+
Device Details
Add Device(s) [Success]
```

```
    Addition of Spine device with ip-address = 10.64.160.30 [Succeeded]
    Addition of Leaf device with ip-address = 10.64.160.38 [Succeeded]
    Addition of Leaf device with ip-address = 10.64.160.11 [Succeeded]
Validate Fabric [Success]
--- Time Elapsed: 10.86306224s -

 efa fabric show -name fs
Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
| 10.64.160.30 |     | SLX-30    | 64512 | Spine | provisioned
| cfg refreshed | DD                | BGP-C,BGPPG-C,INTIP-C, | NA      | 1     |
| 10.64.160.11 |     | SLX-11    | 65000 | Leaf  | provisioned
| cfg refreshed | DD                | MCT-C,MCT-PA,BGP-C,BGPPG-C,INTIP-
C,EVPN-C,O-C,BGPAF-C | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | Leaf  | not
provisioned  | cfg ready | DA                | SYSP-C,MCT-C,MCT-PA,BGP-
C,BGPPG-C,INTIP-C,EVPN-C,O-C,BGPAF-C | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
```

c. Configure the fabric.

   XCO provisions the new leaf with bfd-software-session on its underlay interfaces. Existing devices receive a BGP peer-group refresh to add the new neighbor. Expected result: configure-success, Health: Green. Note: If configure succeeds but BFD sessions cannot be established immediately, XCO reports a warning: 'BFD session(s) could not be established'. This is non-fatal — the configuration is correct and sessions will establish once the remote device completes its BFD negotiation.

```
efa fabric configure  --name fs
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: 48.527742952s ---
 efa fabric show -name fs
Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
| 10.64.160.30 |     | SLX-30    | 64512 | Spine | provisioned
| cfg in-sync | NA                 | NA              | NA      | 1     |
| 10.64.160.11 |     | SLX-11    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                 | NA              | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                 | NA              | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
```

4. Transition an active fabric.

   You can migrate an active hardware BFD fabric to software BFD. Use this step when an existing, fully provisioned fabric is running hardware BFD and must be migrated to software BFD (for example, after an XCO upgrade to 4.1.0).

   a. Confirm current state (hardware BFD, fabric healthy).

      Baseline health check before initiating migration.

      ```
      efa fabric setting show --name fs --advanced | grep BFD | BFD Session Type |
      hardware | efa fabric
      show --name fs # Expected: configure-success, Health: Green, all devices cfg in-sync
      ```

   b. Review pending migration status.

      The detailed fabric view lists the migration type, plane/rack, and the adjacent devices scheduled for BFD-Session-Type migration:

      ```
      efa fabric show --name clos3 --detail

      PENDING MIGRATION STATUS:
      | MIGRATION TYPE    | PLANE | PLANE DEVICE | ROLE   | ADJACENT DEVICE | ROLE |
      | BFD-Session-Type | Spine | 10.20.30.2    | Spine | 10.20.30.10     | Leaf |
      | BFD-Session-Type | MCT    | 10.20.30.10   | Leaf   | 10.20.30.11     | Leaf |
      ```

      Note
      For a non-CLOS fabric, the migration type is Inter-Rack (plus MCT for leaf-internal links), and the output includes the RACK column for each device.

   c. Update BFD session type to software.

      Records migration intent. Fabric status changes to settings-updated. Health changes to Red. No device changes occur at this stage. Updated Fabric Settings indicator: BFD-ST.

      ```
      efa fabric setting update --name "fs" -bfd-session-type software
      fs Fabric Update Successful
      --- Time Elapsed: 4.41886692s ---
      ```

   d. Run the staged migration.

      XCO detects the topology (CLOS ® spine-by-spine). For each Spine in sequence: 1. Admin-down BFD policy on Spine's underlay interfaces 2. Apply bfd-software-session on those interfaces 3. Re-enable BFD policy 4. Verify BFD session re-establishment After all Spines, performs ICL/MCT migration on each Logical Leaf. Expected result: configure-success, Health: Green, all devices cfg in-sync.

      ```
      efa fabric configure  --name fs
      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: 48.527742952s ---
       efa fabric show -name fs
      Fabric Name: fs, 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 |
      +--------------+-----+-----------+-------+-------+-------------+-------------
      +------------------+-----------------+---------+-------+
      | 10.64.160.30 |     | SLX-30    | 64512 | Spine | provisioned
      | cfg in-sync | NA                | NA              | NA      | 1     |
      | 10.64.160.11 |     | SLX-11    | 65000 | Leaf  | provisioned
      | cfg in-sync | NA                | NA              | 2       | 1     |
      ```

```
| 10.64.160.38 |     | SLX-38    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
+-------------+-----+-----------+-------+-------+--------------+-------------
+------------------+-----------------+---------+-------+
```

5. View migration status.

   To view the current migration state, run the following command:

```
efa fabric show --name clos3 --detail

Fabric Name: clos3, Fabric Description: CLOS 3-Stage Production, Fabric Stage: 3,
Fabric Type: clos, Fabric Status: configure-failed, Fabric Health: Black
+-------------+-----+----------+-------+-------+--------------------+-------------
+------------------+-----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  |    DEVICE STATE
|  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+----------+-------+-------+--------------------+-------------
+------------------+-----------------+---------+-------+
| 10.20.30.1  |     | Spine-1   | 64512 | spine | provisioned
| cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.30.2  |     | Spine-2   | 64513 | spine | provisioning failed
| cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.30.10 |     | Leaf-1    | 65000 | leaf  | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.30.11 |     | Leaf-2    | 65000 | leaf  | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.30.12 |     | Leaf-3    | 65001 | leaf  | provisioned
| cfg in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.13 |     | Leaf-4    | 65001 | leaf  | provisioned
| cfg in-sync | NA                | NA              | 3       | 1     |
+-------------+-----+----------+-------+-------+--------------------+-------------
+------------------+-----------------+---------+-------+

PENDING MIGRATION STATUS:
+-----------------+-------+--------------+-------+-----------------+------+
|  MIGRATION TYPE  | PLANE | PLANE DEVICE | ROLE  | ADJACENT DEVICE | ROLE |
+-----------------+-------+--------------+-------+-----------------+------+
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.10     | Leaf |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.11     | Leaf |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.12     | Leaf |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.13     | Leaf |
| BFD-Session-Type | MCT   | 10.20.30.10  | Leaf  | 10.20.30.11     | Leaf |
| BFD-Session-Type | MCT   | 10.20.30.12  | Leaf  | 10.20.30.13     | Leaf |
+-----------------+-------+--------------+-------+-----------------+------+

efa fabric show --name clos5 --detail

Fabric Name: clos5, Fabric Description: CLOS 5-Stage DC, Fabric Stage:
5, Fabric Type: clos, Fabric Status: configure-failed, Fabric Health: Black
+-------------+-----+--------------+-------+-------------+---------------------
+-------------+------------------+-----------------+---------+-------+
| IP ADDRESS  | POD |  HOST NAME   |  ASN  |    ROLE     |    DEVICE
STATE     |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+--------------+-------+-------------+---------------------
+-------------+------------------+-----------------+---------+-------+
| 10.20.30.1  |     | Spine-1      | 64512 | spine       | provisioning
failed | cfg in-sync | NA               | NA              | NA      | 1     |
| 10.20.30.2  |     | Spine-2      | 64513 | spine       | provisioned
| cfg in-sync | NA               | NA              | NA      | 1     |
| 10.20.30.5  |     | SuperSpine-1 | 64520 | super-spine | provisioned
| cfg in-sync | NA               | NA              | NA      | 1     |
| 10.20.30.6  |     | SuperSpine-2 | 64521 | super-spine | provisioned
| cfg in-sync | NA               | NA              | NA      | 1     |
| 10.20.30.10 |     | Leaf-1       | 65000 | leaf        | provisioned
| cfg in-sync | NA               | NA              | 2       | 1     |
```

```
| 10.20.30.11 |     | Leaf-2       | 65000 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.30.12 |     | Leaf-3       | 65001 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.13 |     | Leaf-4       | 65001 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.14 |     | BorderLeaf-1 | 65010 | border-leaf | provisioned
| cfg in-sync | NA                | NA              | 4       | 1     |
| 10.20.30.15 |     | BorderLeaf-2 | 65010 | border-leaf | provisioned
| cfg in-sync | NA                | NA              | 4       | 1     |
+-------------+-----+--------------+-------+-------------+----------------------
+-------------+-------------------+-----------------+---------+-------+

PENDING MIGRATION STATUS:
+------------------+-------+--------------+------------+-----------------+------------+
|  MIGRATION TYPE  | PLANE | PLANE DEVICE |    ROLE    | ADJACENT DEVICE |    ROLE    |
+------------------+-------+--------------+------------+-----------------+------------+
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.5      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.6      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.10     | Leaf       |
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.11     | Leaf       |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.5      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.6      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.12     | Leaf       |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.13     | Leaf       |
| BFD-Session-Type | MCT   | 10.20.30.10  | Leaf       | 10.20.30.11     | Leaf       |
| BFD-Session-Type | MCT   | 10.20.30.12  | Leaf       | 10.20.30.13     | Leaf       |
| BFD-Session-Type | MCT   | 10.20.30.14  | BorderLeaf | 10.20.30.15     | BorderLeaf |
+------------------+-------+--------------+------------+-----------------+------------+
```

efa fabric show --name nonclos-2rack --detail

```
Fabric Name: nonclos-2rack, Fabric Description: Non-CLOS 2 Racks, Fabric Stage: NA,
Fabric Type: non-clos, Fabric Status: configure-failed, Fabric Health: Black
+-------------+--------+-----------+-------+------+---------------------+-------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  |  RACK  | HOST NAME |  ASN  | ROLE |    DEVICE STATE     |  APP
STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+--------+-----------+-------+------+---------------------+-------------
+-------------------+-----------------+---------+-------+
| 10.20.30.10 | Rack-A | Leaf-1    | 65000 | leaf | provisioned         | cfg in-
sync | NA                | NA              | 2       | 1     |
| 10.20.30.11 | Rack-A | Leaf-2    | 65000 | leaf | provisioned         | cfg in-
sync | NA                | NA              | 2       | 1     |
| 10.20.30.20 | Rack-B | Leaf-3    | 65001 | leaf | provisioning failed | cfg
in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.21 | Rack-B | Leaf-4    | 65001 | leaf | provisioned         | cfg in-
sync | NA                | NA              | 3       | 1     |
+-------------+--------+-----------+-------+------+---------------------+-------------
+-------------------+-----------------+---------+-------+

PENDING MIGRATION STATUS:
+------------------+------------+--------------+------+--------+-----------------
+------+--------+
|  MIGRATION TYPE  |   PLANE    | PLANE DEVICE | ROLE |  RACK
| ADJACENT DEVICE | ROLE |  RACK  |
+------------------+------------+--------------+------+--------+-----------------
+------+--------+
| BFD-Session-Type | Inter-Rack | 10.20.30.20  | Leaf | Rack-B
| 10.20.30.10     | Leaf | Rack-A |
| BFD-Session-Type | Inter-Rack | 10.20.30.20  | Leaf | Rack-B
| 10.20.30.11     | Leaf | Rack-A |
| BFD-Session-Type | Inter-Rack | 10.20.30.21  | Leaf | Rack-B
| 10.20.30.10     | Leaf | Rack-A |
| BFD-Session-Type | Inter-Rack | 10.20.30.21  | Leaf | Rack-B
```

```
| 10.20.30.11     | Leaf | Rack-A |
| BFD-Session-Type | MCT        | 10.20.30.20  | Leaf | Rack-B
| 10.20.30.21     | Leaf | Rack-B |
+------------------+------------+--------------+------+--------+-----------------
+------+--------+
```

The output displays the **Pending Migration Status** section, including:

- Migration type
- Spine or Rack being migrated
- Affected devices
- Adjacent devices
- Migration scope information

6. Update BFD settings, add devices, and then configure.

   It covers the scenario where a BFD setting update and a device addition are both pending before the next configure operation. XCO handles both changes in a single configure run.

   a. Update BFD session type on active fabric.

      Fabric status: settings-updated, Health: Red. Pending setting: BFD-ST.

      ```
      efa fabric setting update --name "fs" -bfd-session-type software
      fs Fabric Update Successful
      --- Time Elapsed: 4.41886692s ---
      ```

   b. Add a new leaf device.

      Fabric health transitions to Black (settings-updated + new device pending). Updated Fabric Settings still shows BFD-ST.

      ```
      efa fabric device add --name "fs" --ip 10.64.160.38 --role leaf --hostname
      SLX-160-38
      Inventory Device(s) Registration[success]
      +----+--------------+-----------+-------+--------------+---------------+
      | ID |  IP Address  | Host Name | Model | Chassis Name |   Firmware    |
      +----+--------------+-----------+-------+--------------+---------------+
      | 5  | 10.64.160.38 | SLX-38    | 3012  | Extreme 8720-32C | 20.8.1_po_bfd |
      +----+--------------+-----------+-------+--------------+---------------+
      Device Details
      Add Device(s) [Success]

          Addition of Spine device with ip-address = 10.64.160.30 [Succeeded]
          Addition of Leaf device with ip-address = 10.64.160.38 [Succeeded]
          Addition of Leaf device with ip-address = 10.64.160.11 [Succeeded]
      Validate Fabric [Success]
      --- Time Elapsed: 10.86306224s -
      ```

   c. Run combined configure.

      XCO handles both changes simultaneously: - Performs spine-by-spine software BFD migration on existing devices. - Provisions the new leaf device directly with bfd-software-session. Expected result: configure-success, Health: Green, all three devices cfg in-sync.

      ```
      efa fabric configure  --name fs
      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: 48.527742952s ---
       efa fabric show -name fs
      Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+-------------+------------
+------------------+----------------+---------+-------+
| 10.64.160.30 |     | SLX-30    | 64512 | Spine | provisioned
| cfg in-sync | NA                | NA             | NA      | 1     |
| 10.64.160.11 |     | SLX-11    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                | NA             | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | Leaf  | provisioned
| cfg in-sync | NA                | NA             | 2       | 1     |
+--------------+-----+-----------+-------+-------+-------------+------------
+------------------+----------------+---------+-------+
```

7. Run the **`efa fabric show`** command for Clos and non-Clos fabric.

Show output for Clos fabric

```
efa fabric show --name clos3 --detail

Fabric Name: clos3, Fabric Description: CLOS 3-Stage Production, Fabric Stage: 3,
Fabric Type: clos, Fabric Status: configure-failed, Fabric Health: Black
+-------------+-----+-----------+-------+-------+---------------------+-------------
+------------------+----------------+---------+-------+
| IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE  |    DEVICE STATE
|  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+-----------+-------+-------+---------------------+-------------
+------------------+----------------+---------+-------+
| 10.20.30.1  |     | Spine-1   | 64512 | spine | provisioned
| cfg in-sync | NA                | NA             | NA      | 1     |
| 10.20.30.2  |     | Spine-2   | 64513 | spine | provisioning failed
| cfg in-sync | NA                | NA             | NA      | 1     |
| 10.20.30.10 |     | Leaf-1    | 65000 | leaf  | provisioned
| cfg in-sync | NA                | NA             | 2       | 1     |
| 10.20.30.11 |     | Leaf-2    | 65000 | leaf  | provisioned
| cfg in-sync | NA                | NA             | 2       | 1     |
| 10.20.30.12 |     | Leaf-3    | 65001 | leaf  | provisioned
| cfg in-sync | NA                | NA             | 3       | 1     |
| 10.20.30.13 |     | Leaf-4    | 65001 | leaf  | provisioned
| cfg in-sync | NA                | NA             | 3       | 1     |
+-------------+-----+-----------+-------+-------+---------------------+-------------
+------------------+----------------+---------+-------+

PENDING MIGRATION STATUS:
+------------------+-------+--------------+-------+-----------------+------+
|  MIGRATION TYPE  | PLANE | PLANE DEVICE | ROLE  | ADJACENT DEVICE | ROLE |
+------------------+-------+--------------+-------+-----------------+------+
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.10     | Leaf |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.11     | Leaf |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.12     | Leaf |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine | 10.20.30.13     | Leaf |
| BFD-Session-Type | MCT   | 10.20.30.10  | Leaf  | 10.20.30.11     | Leaf |
| BFD-Session-Type | MCT   | 10.20.30.12  | Leaf  | 10.20.30.13     | Leaf |
+------------------+-------+--------------+-------+-----------------+------+

efa fabric show --name clos5 --detail

Fabric Name: clos5, Fabric Description: CLOS 5-Stage DC, Fabric Stage:
5, Fabric Type: clos, Fabric Status: configure-failed, Fabric Health: Black
+-------------+-----+--------------+-------+-------------+---------+---------------------
+-------------+-------------------+-----------------+---------+-------+
| IP ADDRESS  | POD |  HOST NAME   |  ASN  |    ROLE     |    DEVICE
STATE     |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-----+--------------+-------+-------------+---------------------
+-------------+-------------------+-----------------+---------+-------+
| 10.20.30.1  |     | Spine-1      | 64512 | spine       | provisioning
```

```
failed | cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.30.2  |     | Spine-2      | 64513 | spine       | provisioned
| cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.30.5  |     | SuperSpine-1 | 64520 | super-spine | provisioned
| cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.30.6  |     | SuperSpine-2 | 64521 | super-spine | provisioned
| cfg in-sync | NA                | NA              | NA      | 1     |
| 10.20.30.10 |     | Leaf-1       | 65000 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.30.11 |     | Leaf-2       | 65000 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.30.12 |     | Leaf-3       | 65001 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.13 |     | Leaf-4       | 65001 | leaf        | provisioned
| cfg in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.14 |     | BorderLeaf-1 | 65010 | border-leaf | provisioned
| cfg in-sync | NA                | NA              | 4       | 1     |
| 10.20.30.15 |     | BorderLeaf-2 | 65010 | border-leaf | provisioned
| cfg in-sync | NA                | NA              | 4       | 1     |
+-------------+-----+--------------+-------+-------------+---------------------
+-------------+-------------------+-----------------+---------+-------+

PENDING MIGRATION STATUS:
+------------------+-------+--------------+------------+-----------------+------------+
|  MIGRATION TYPE  | PLANE | PLANE DEVICE |    ROLE    | ADJACENT DEVICE |    ROLE    |
+------------------+-------+--------------+------------+-----------------+------------+
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.5      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.6      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.10     | Leaf       |
| BFD-Session-Type | Spine | 10.20.30.1   | Spine      | 10.20.30.11     | Leaf       |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.5      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.6      | SuperSpine |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.12     | Leaf       |
| BFD-Session-Type | Spine | 10.20.30.2   | Spine      | 10.20.30.13     | Leaf       |
| BFD-Session-Type | MCT   | 10.20.30.10  | Leaf       | 10.20.30.11     | Leaf       |
| BFD-Session-Type | MCT   | 10.20.30.12  | Leaf       | 10.20.30.13     | Leaf       |
| BFD-Session-Type | MCT   | 10.20.30.14  | BorderLeaf | 10.20.30.15     | BorderLeaf |
+------------------+-------+--------------+------------+-----------------+------------+
```

Show output for non-Clos fabric

```
efa fabric show --name nonclos-2rack --detail

Fabric Name: nonclos-2rack, Fabric Description: Non-CLOS 2 Racks, Fabric Stage: NA,
Fabric Type: non-clos, Fabric Status: configure-failed, Fabric Health: Black
+-------------+--------+-----------+-------+------+---------------------+-------------
+-------------------+-----------------+---------+-------+
| IP ADDRESS  |  RACK  | HOST NAME |  ASN  | ROLE |    DEVICE STATE     |  APP
STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-------------+--------+-----------+-------+------+---------------------+-------------
+-------------------+-----------------+---------+-------+
| 10.20.30.10 | Rack-A | Leaf-1    | 65000 | leaf | provisioned         | cfg in-
sync | NA                | NA              | 2       | 1     |
| 10.20.30.11 | Rack-A | Leaf-2    | 65000 | leaf | provisioned         | cfg in-
sync | NA                | NA              | 2       | 1     |
| 10.20.30.20 | Rack-B | Leaf-3    | 65001 | leaf | provisioning failed | cfg
in-sync | NA                | NA              | 3       | 1     |
| 10.20.30.21 | Rack-B | Leaf-4    | 65001 | leaf | provisioned         | cfg in-
sync | NA                | NA              | 3       | 1     |
+-------------+--------+-----------+-------+------+---------------------+-------------
+-------------------+-----------------+---------+-------+

PENDING MIGRATION STATUS:
+------------------+------------+--------------+------+--------+-----------------
+------+--------+
|  MIGRATION TYPE  |   PLANE    | PLANE DEVICE | ROLE |  RACK
```

```
| ADJACENT DEVICE | ROLE |  RACK  |
+-----------------+-----------+-------------+------+--------+----------------
+------+--------+
| BFD-Session-Type | Inter-Rack | 10.20.30.20  | Leaf | Rack-B
| 10.20.30.10     | Leaf | Rack-A |
| BFD-Session-Type | Inter-Rack | 10.20.30.20  | Leaf | Rack-B
| 10.20.30.11     | Leaf | Rack-A |
| BFD-Session-Type | Inter-Rack | 10.20.30.21  | Leaf | Rack-B
| 10.20.30.10     | Leaf | Rack-A |
| BFD-Session-Type | Inter-Rack | 10.20.30.21  | Leaf | Rack-B
| 10.20.30.11     | Leaf | Rack-A |
| BFD-Session-Type | MCT        | 10.20.30.20  | Leaf | Rack-B
| 10.20.30.21     | Leaf | Rack-B |
+-----------------+-----------+-------------+------+--------+----------------
+------+--------+
```

Compliance, Security, and Miscellaneous

Topic	Assessment
ISO 14001 — Environmental Management	Not applicable. This feature introduces no hardware components, manufacturing processes, or materials with environmental impact considerations.
GDPR — Data Privacy	Not applicable. This feature does not collect, process, store, or transmit personally identifiable information (PII). All operational state (BFD session counters, adjacency status) is transient, non-personal, and automatically cleared on device reboot.
Third-Party and Open-Source Dependencies	None. This feature is implemented entirely within the XCO and SLX-OS software stacks using existing BFD and BGP protocol implementations. No new third-party packages are introduced.
Security	This feature does not implement encryption, authentication, or key management. It operates within the existing secure management channel (NETCONF over SSH) used by XCO for all device provisioning. No new attack surface is introduced.
Checkpointing and Backward Compatibility	Fabric configurations created with XCO 4.0.x (hardware BFD) continue to operate unchanged after upgrading to XCO 4.1.0. The new BFD session type field is backward-compatible and defaults to hardware for all pre-existing fabrics.
New APIs	No new REST or gRPC API endpoints are introduced. The existing `fabric setting update` and `fabric show-config` APIs are extended with the `bfd_session_type` field.

Topic	Assessment
Telegraf / Cloud Management APIs	No new Telegraf metrics or cloud management API messages are introduced for this feature.
Universal Port Management (UPM)	This feature has no interaction with UPM. BFD session-type configuration is applied exclusively to XCO-managed fabric underlay interfaces.
Licensing	No additional license is required. BFD session-type selection is included within the standard IP Fabric license entitlement.

Fabric Post-Upgrade Behavior

Learn about the post-upgrade behavior of a fabric.

Upgrade Success Case

The fabric retains its pre-upgrade status, but device state might change based on the comparison between asset and fabric details.

Before Upgrade

Shows fabric status as `configure-success` and health as `Green`

```
efa fabric show --name fs --detail

Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.160.11 |     | SLX-11    | 64512 | spine |
provisioned  | cfg in-sync | NA                | NA              | NA      | 1     |
| 10.64.160.30 |     | SLX-30    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

Fabric Health

```
Fabric Level Physical Topology Health     :  Green
-------------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                              :  10.64.160.11 [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]                               :  10.64.160.30 [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]                              :  10.64.160.38 [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
-------------------------------------------------------------------------------

===============================================================================
CONFIG GEN REASON:
```

.

After Upgrade

Shows no changes in fabric status and health.

```
efa fabric show --name fs --detail

Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.160.11 |     | SLX-11    | 64512 | spine |
provisioned  | cfg in-sync | NA                | NA              | NA      | 1     |
| 10.64.160.30 |     | SLX-30    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

Fabric Health

```
 Fabric Level Physical Topology Health     :  Green
-------------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                               :  10.64.160.11 [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]                               :  10.64.160.30 [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]                               :  10.64.160.38 [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
-------------------------------------------------------------------------------

===============================================================================

CONFIG GEN REASON:
```

Upgrade Success Failure Case

Shows fabric status as `upgrade-failed` and health as `Red`, with device upgrade status indicating failures. If upgrade fails fabric health must degrade.

```
efa fabric show --name fs --detail

Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric Status:
upgrade-failed, Fabric Health: Red
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+----------------+---------+-------+
|  IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE
| DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+----------------+---------+-------+
| 10.64.160.11 |     | SLX-11    | 64512 | spine
| provisioned  | cfg in-sync | NA               | NA              | NA      | 1     |
| 10.64.160.30 |     | SLX-30    | 65000 | leaf
| provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | leaf
| provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+----------------+---------+-------+
```

Device Upgrade Status

```
+--------------+--------------------+
|  IP ADDRESS  |   UPGRADE STATUS   |
+--------------+--------------------+
| 10.64.160.31 | LALD-UpgradeFailed |
| 10.64.160.30 | LALD-UpgradeFailed |
| 10.64.160.38 | LALD-UpgradeFailed |
+--------------+--------------------+
```

Fabric Health

```
 Fabric Level Physical Topology Health     :  Green
-------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                                :  10.64.160.11 [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]                                :  10.64.160.30 [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]                                :  10.64.160.38 [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
-------------------------------------------------------------------------------

===============================================================================

CONFIG GEN REASON:
```

Upgrade Recovery After Failure Case

Learn about the upgrade recovery (full and partial) after the failure cases.

Case 1: Full Recovery

All devices are in-sync (`cfg in-sync`) and all the fabric tables are in none config (no pending configs). The **`efa fabric show`** command output shows fabric status as `configure-success` and health as `Green`.

```
efa fabric show --name fs
Fabric Name: fs, 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 |
+--------------+-----+-----------+-------+-------+--------------+-------------
+------------------+----------------+---------+-------+
| 10.64.160.11 |     | SLX-11    | 64512 | spine |
provisioned  | cfg in-sync | NA               | NA              | NA      | 1     |
| 10.64.160.30 |     | SLX-30    | 65000 | leaf  |
provisioned  | cfg in-sync | NA               | NA              | 2       | 1     |
```

```
| 10.64.160.38 |     | SLX-38    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

Case 2: Partial Recovery

After upgrade recovery, some devices are still in `cfg-refreshed` state, needing a config refresh. The **`efa fabric show`** command output shows fabric status as `upgrade-success` and health as `Black`, with affected devices indicating `cfg refreshed` state.

```
efa fabric show --name fs
Fabric Name: fs, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: upgrade-success, Fabric Health: Black
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
|  IP ADDRESS  | POD | HOST NAME |  ASN  | ROLE
| DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.160.11 |     | SLX-11    | 64512 | spine | provisioned
| cfg refreshed | BGPU                | BGPU              | NA      | 1     |
| 10.64.160.30 |     | SLX-30    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.64.160.38 |     | SLX-38    | 65000 | leaf  |
provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+-----+-----------+-------+-------+--------------+-------------
+-------------------+-----------------+---------+-------+
```

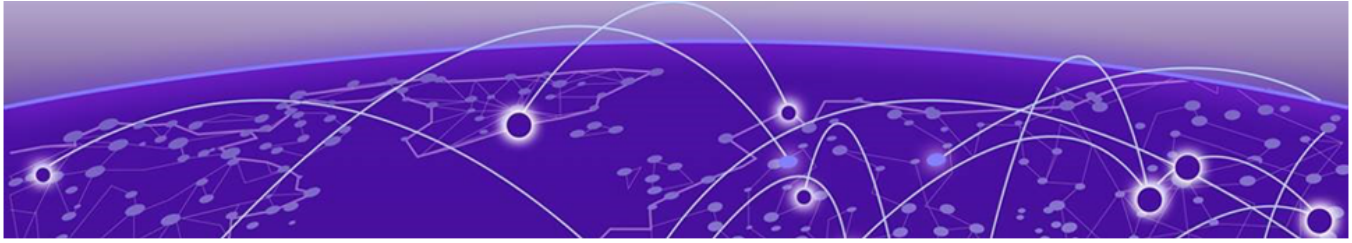

Tenant Service Provisioning

Learn about configuring tenants, tenant networks, endpoints, and traffic mirror session.

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.

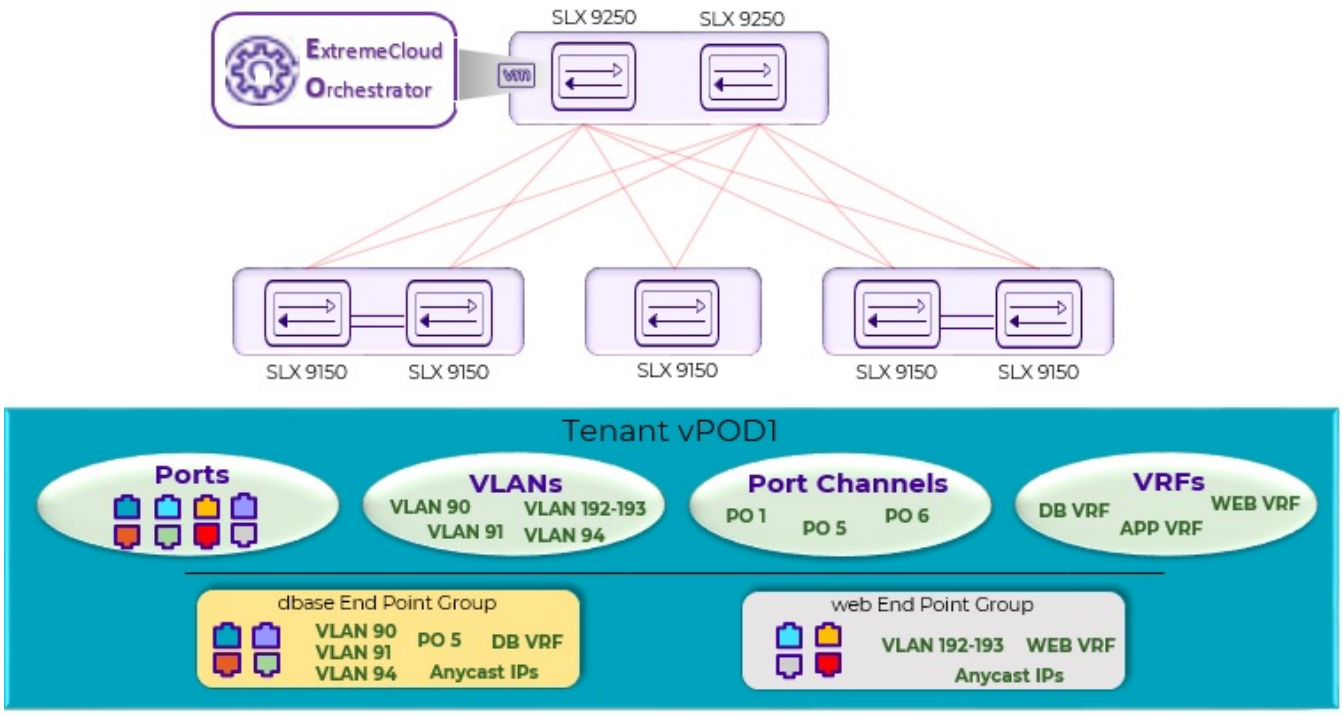

Figure 16: Tenant Services Overview

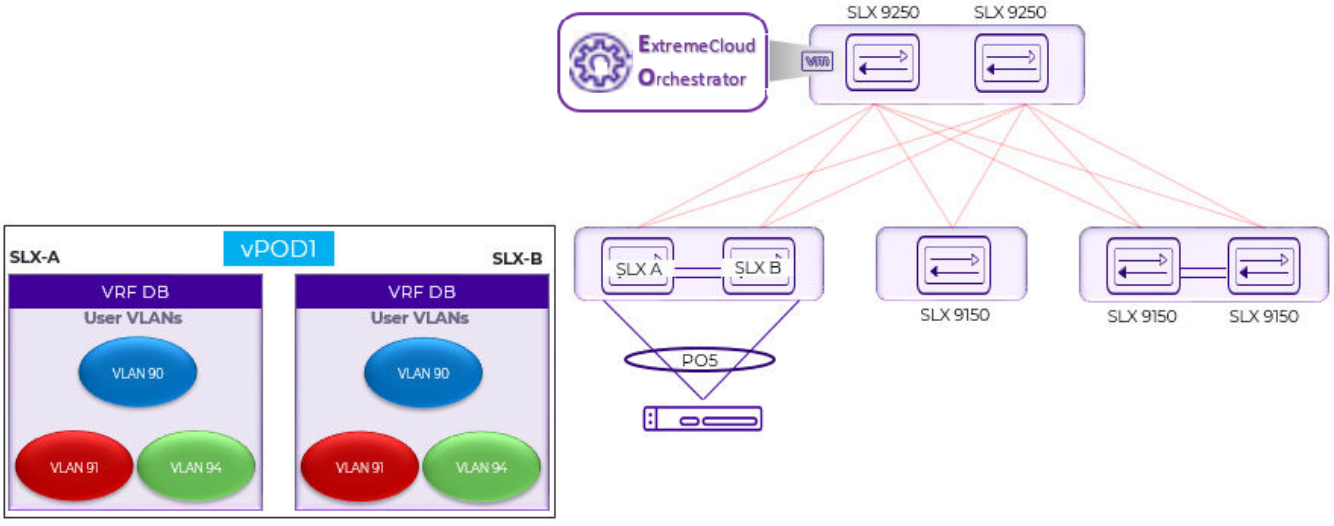

Figure 17: Tenant Name vPOD1, VRF Name DB

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.

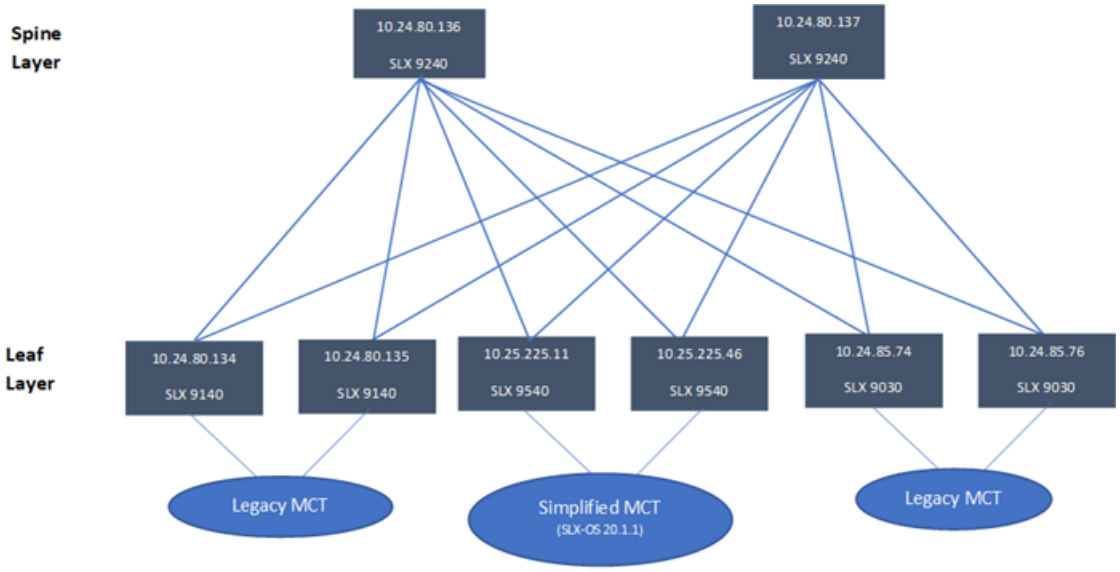

Figure 18: 3-stage Clos topology

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 <username> --password
<password>
--name fabric1

efa fabric configure --name fabric1
```

The following figure shows tenant constructs in the Clos Fabric.

Figure 19: 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 286
2. Update a Tenant on page 287
3. Show a Tenant on page 289
4. Delete a Tenant on page 291
5. Configure a Tenant on page 291
6. Share Resources Across Tenants using Shared Tenant on page 551

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, 4.1.0*.

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.

Note
For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

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, 4.1.0*.

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 292
2. Create a Shared Tenant on page 293
3. Scalability on page 294
4. VLAN-based Tenant on page 294
5. Bridge Domain-based Tenant on page 294
6. OS ONE VLAN-based Tenant on page 295

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, 4.1.0*.

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, 4.1.0*.

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 ---
```

OS ONE VLAN-based Tenant

Use this topic to learn about design limitations in the existing VLAN-based EPG model (v4.0) by introducing unified handling of Bridge Domains (BDs) and VLAN-based BDs across tenants, inventory, and device provisioning workflows.

Key Changes

- Unified BD and VLAN-based BD handling
- Inventory REST API changes
- Drift reconcile CLI output changes
- VNI allocation changes in homogeneous and heterogeneous fabrics

Inventory REST API Changes

- `GET /inventory/bridgedomain`: Returns BDs with LIFs and VLAN IDs

```
http://localhost:8082/v1/inventory/bridgedomain?device_ip=10.64.188.68
[
    {
        "bd": {
            "id": 37,
            "name": "1",
            "type": "p2mp"
        },
        "lifs": [
            {
                "id": 91,
                "lif_name": "0/3.13",
                "name": "0/3.13",
                "vlan_id": 13,
                "vlan_tag": "tagged",
                "interface": {
                    "id": 145,
                    "int_type": "ethernet",
                    "name": "0/3"
                }
            },
            {
                "id": 93,
                "lif_name": "0/4.14",
                "name": "0/4.14",
                "vlan_id": 14,
                "vlan_tag": "tagged",
                "interface": {
                    "id": 147,
                    "int_type": "ethernet",
                    "name": "0/4"
                }
            },
          ]
    },
    {
        "bd": {
            "id": 39,
            "name": "2",
            "type": "vlan",
            "vlan_id": 20
        },
        "lifs": [
            {
```

```
                "id": 97,
                "lif_name": "0/10.20",
                "name": "0/10.20",
                "vlan_id": 20,
                "vlan_tag": "tagged",
                "interface": {
                    "id": 159,
                    "int_type": "ethernet",
                    "name": "0/10"
                }
            },
            {
                "id": 99,
                "lif_name": "0/11.20",
                "name": "0/11.20",
                "vlan_id": 20,
                "vlan_tag": "tagged",
                "interface": {
                    "id": 161,
                    "int_type": "ethernet",
                    "name": "0/11"
                }
            }
        ]
    }

]
```

- `GET /inventory/vlans`: Returns empty response after redesign

```
http://localhost:8082/v1/inventory/vlans?device_ip=10.64.188.68
  {}
```

Drift Reconcile CLI Output Changes

- BD drift changes: Shows BD ID and App-state

```
=============== BD Drift ================================
BD         : 2
Vlan-id    : 5
App-state  : cfg-refreshed

=============== EVPN Drift ================================
EVPN            : ncls-fab-ave
App-state       : cfg-in-sync

Drifted BD's
+----+---------------+
| BD |   App-State   |
+----+---------------+
| 2  | cfg-refreshed |
-------+----------------
```

- EVPN drift changes: Shows EVPN name and App-state

```
# show running-config bridge-domain
bridge-domain 1 mode vlan
description Tenant Default Native VLAN
vlan-id 1
member ethernet 0/5 untagged
!
bridge-domain 2 mode vlan
description Tenant L3 Extended VLAN
vlan-id 5
member ethernet 0/5
arp-proxy
suppress-arp
```

```
!
!
bridge-domain 3 mode vlan
description Tenant L3 Extended VLAN
vlan-id 6
member ethernet 0/5
arp-proxy
suppress-arp
!
!

If BD for "ctag: 5 i.e bridge-domain 2" configuration is deleted,
```

- MCT drift changes: Shows MCT cluster BD mapping

```
============== MCT Cluster Drift ================================
Cluster         : fabpktos-cluster-1
Cluster ID      : 19

Drifted Mct cluster BD Mapping
+------+---------------+
| BD   |   App-State   |
+------+---------------+
| 1    | cfg-refreshed |
+------+---------------+
```

Limitations

- The `reserve-first-4k-vni-for-vlan` flag cannot be updated after EPGs exist.
- SLX device addition not allowed if `reserve-first-4k-vni-for-vlan` is NO.

VNI Allocation

Use this topic to learn about the VNI allocation in SLX OS, OS ONE, homogeneous, and Heterogeneous fabric.

- SLX OS fabric: Allocates L2VNI based on Ctag (default)
- OS ONE fabric: Allocates lowest available BD across fabric
- Heterogeneous fabric: Supports both SLX OS and OS ONE allocation logic

VNI Allocation in Homogeneous Fabric

Fabric Setting Flag

- `--reserve-first-4k-vni-for-vlan`: Controls VNI allocation for VLAN-based tenants
- Default value: Same as `--vni_auto_map` for both SLX and OS ONE fabrics

Restrictions

- Cannot update flag after EPGs exist
- Cannot set flag to Yes when vni-auto-map is No
- Cannot set flag to No when SLX device already exists

Behavior

- Updating `--vni-auto-map` updates `--reserve-first-4k-vni-for-vlan` (with warning)
- Updating `--reserve-first-4k-vni-for-vlan` does not update `--vni-auto-map`

VNI Allocation Examples

SLX OS Fabric (`--reserve-first-4k-vni-for-vlan` = true)

- VLAN-based tenant: L2VNI = Ctag (for example, 415)
- BD-based tenant: L2VNI = BD number + 4096 (for example, 4097)

OS ONE Fabric (`--reserve-first-4k-vni-for-vlan` = false)

- VLAN-based tenant: L2VNI = Lowest available BD (for example, 5)
- BD-based tenant: L2VNI = Lowest available BD (for example, 3)

VNI Allocation in Heterogeneous Fabric (Brownfield)

Brownfield Deployment

- Existing SLX OS fabric transitions to hybrid fabric with `--reserve-first-4k-vni-for-vlan` set to true. `--reserve-first-4k-vni-for-vlan` defaults to true.
- OS ONE devices added to fabric
- VNI allocation
  - OS ONE: L2VNI = Lowest available BD (for example, 3, 5)
  - SLX: L2VNI = BD number + 4096 (for example, 4097, 4101)

Greenfield Deployment

- Creating OS ONE fabric without SLX OS devices:
  - `--reserve-first-4k-vni-for-vlan` set to true by default

```
+------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------|-------|------------+
IP ADDRESS   | RACK | HOST NAME | ASN        | ROLE | DEVICE STATE | APP STATE
| CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID | MPLS LB ID |
10.64.184.83 | r1   | vtos83    | 4200000000 | Leaf | provisioned  | cfg in-sync
| NA                | NA              | 2       | 1     |            |
10.64.184.82 | r1   | vtos82    | 4200000000 | Leaf | provisioned  | cfg in-sync
| NA                | NA              | 2       | 1     |            |
+------------+------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+------------+
# efa fabric setting update --name fabrichyd --reserve-first-4k-vni-for-vlan No
fabrichyd Fabric Update Successful
```

Adding SLX device

```
efa fabric device add -name fabrichyd  --ip 10.64.64.44 --rack "r1"
Error: SLX device addition is not allowed as  reserve-first-4k-vni-for-vlan is
NO

# efa tenant epg show for bd network
+------+------+-------+---------+--------------+--------------+----------
+----------+---------------+--------+---------+---------+---------+-----------
+------------+
| 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 | | | +------+---------------------------+-------+---------
+--------------+--------------+----------+-----------------------+---------------
+--------+---------+---------------+--------------+------------
| 415  | Tenant L2 Extended BD |   3   | Auto-BD-3
|              |              |   F/F    |                       |      F/F
|        |         |     false      |    false     | provisioned | cfg-in-sync |
+------+-----------------------+-------+----------+--------------+--------------
+----------+----------------------+---------------+--------+---------
```

```
+----------------+-------------+------------+-------------+


Efa tenant epg show for vlan network
+------+----------------------------+-------+---------+--------------
+--------------+----------+----------------------+------------
| 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 | | | +------+----------------------------+-------+---------
+--------------+-------------+----------+----------------------+---------------
+--------+---------+----------------+-------------+-------------+-------------+
|  1   | Tenant Default Native VLAN |   4   |
|              |              |   F/F    |                        |      F/F
|        |         |     false      |    false     | provisioned | cfg-in-sync |
+------+----------------------------+-------+---------+--------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
| 415  |  Tenant L2 Extended VLAN   |   5   |
|              |              |   F/F    |                        |      F/F
|        |         |     false      |    false     | provisioned | cfg-in-sync |
+------+----------------------------+-------+---------+--------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
Network Property [Flags : * - Native Vlan]
```

  - Can be changed to false if no SLX OS devices will be added

- Adding SLX OS devices:
  - Requires `--reserve-first-4k-vni-for-vlan` set to true
  - VNI allocation follows SLXOS logic (for example, L2VNI = Ctag or BD number + 4096)

```
+--------------+------+-----------+------------+------+--------------
+-------------+------------------+-----------------+---------
| IP ADDRESS | RACK | HOST NAME | ASN | ROLE | DEVICE STATE |
APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID | MPLS LB ID |
10.64.184.83 | r1 | vtos83 | 4200000000 |
Leaf | provisioned | cfg in-sync | NA | NA | 2 | 1 | | |
10.64.184.82 | r1 | vtos82 | 4200000000 |
Leaf | provisioned | cfg in-sync | NA | NA | 2 | 1 | |
+--------------+------+-----------+------------+------+--------------
+-------------+------------------+-----------------+---------
# efa fabric setting update --name fabrichyd --reserve-first-4k-vni-for-vlan Yes
fabrichyd Fabric Update Successful

Adding SLX device
efa fabric device add -name fabrichyd  --ip 10.64.64.44 --rack "r1"
+--------------+------+-----------+------------+------+--------------
+-------------+------------------+-----------------+---------
| IP ADDRESS | RACK | HOST NAME | ASN | ROLE | DEVICE STATE |
APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID | MPLS LB ID |
10.64.184.83 | r1 | vtos83 | 4200000000 |
Leaf | provisioned | cfg in-sync | NA | NA | 2 | 1 | | |
10.64.184.82 | r1 | vtos82 | 4200000000 |
Leaf | provisioned | cfg in-sync | NA | NA | 2 | 1 | |
10.64.164.44 | r1 | SLX1 | 4200000000 |
Leaf | provisioned | cfg in-sync | NA | NA | 2 | 1 | |
+--------------+------+-----------+------------+------+--------------
+-------------+------------------+-----------------+---------

Efa tenant epg show for vlan network
+------+----------------------------+-------+---------+--------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
```

```
| 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 |             |             |
+------+-----------------------------+-------+---------+-------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
|  1   | Tenant Default Native VLAN |   1   |
|              |              |   F/F    |                       |      F/F
|        |         |     false      |    false     | provisioned | cfg-in-sync |
+------+-----------------------------+-------+---------+-------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
| 415  |  Tenant L2 Extended VLAN   |  415  |
|              |              |   F/F    |                       |      F/F
|        |         |     false      |    false     | provisioned | cfg-in-sync |




Efa tenant epg show for BDnetwork

+------+-----------------------+-------+-------------+-------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
| 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 |             |             |
+------+-----------------------+-------+-------------+-------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
| 415  | Tenant L2 Extended BD | 4097  | Auto-BD-4097
|              |              |   F/F    |                       |      F/F
|        |         |     false      |    false     | provisioned | cfg-in-sync |
+------+-----------------------+-------+-------------+-------------
+--------------+----------+----------------------+---------------+--------
+---------+----------------+-------------+-------------+-------------+
```

SwitchPortMode Trunk Attributes and EPG Operations

Use this topic to learn about the supported SwitchPortMode trund attributes and EPG operations.

Supported SwitchPortMode Trunk Attributes on OS ONE

VLAN EPG	BD EPG
• switchport-mode trunk • switchport-mode trunk + switchport-native-vlan-tagging • switchport-mode trunk + switchport-native-vlan • switchport-mode trunk + switchport-native-vlan + switchport-native-vlan-tagging • switchport-mode trunk-no-default-native	• switchport-mode trunk • switchport-mode trunk + switchport-native-vlan-tagging • switchport-mode trunk-no-default-native

EPG Update Operations with Port-Group-Add (OS ONE to SLX-OS)

Supported Operations	Non-Supported Operations
• Ctag-range-add (BD) • Ctag-range-delete (VLAN) • Vrf-add • Vrf-delete • Anycast-ip-add • Anycast-ip-delete • Local-ip-add • Local-ip-delete • Suppress-arp (BD) • Suppress-nd (BD)	• Ctag-range-add (VLAN) • Vrf-delete (with dependencies) • Anycast-ip-delete (with dependencies)

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 301
2. Update a Port Channel on page 304
3. Delete a Port Channel on page 306
4. Show a Port Channel on page 308
5. Configure a Port Channel on page 310

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, 4.1.0*.

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.

Note

For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

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 Port Channel, VRF, or EPG delete operations, use a single-name format for all entities within a tenant.
- For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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 ---
```

Delete Pending Port Channel Configuration

You can delete pending configuration on a port channel.

About This Task

Follow this procedure to push or remove the pending configuration on a port channel.

Procedure

Run the following command:

```
efa tenant po configure
```

The **`efa tenant po configure`** command pushes or removes the pending configuration of a port channel when it is in one of the following states:

`po-delete-pending | po-port-delete-pending | po-lacp-timeout-set-pending | po-description-set-pending | po-min-links-count-set-pending | po-mtu-delete-pending`

Example

```
efa tenant po show
+------+--------+----+--------+------+-------------+----------+---------+-----------------+------------+-------------
+-----------------------+
| Name | Tenant | ID | Speed  | MTU  | Negotiation | Min Link |  Lacp   |      Ports      |   State    |
Dev State  |        App State      |
|      |        |    |        |      |             |  Count   | Timeout |                 |            |
|                    |
+------+--------+----+--------+------+-------------+----------+---------+-----------------+------------+-------------
+-----------------------+
| po1  |  tv3   | 1  | 10Gbps | 9216 |   active    |    1     |  short  | 10.20.61.91[0/4] | po-created
| provisioned | po-mtu-delete-pending |
|      |        |    |        |      |             |          |         | 10.20.61.90[0/4] |            |
|                    |
+------+--------+----+--------+------+-------------+----------+---------+-----------------+------------+-------------
+-----------------------+
Port Channel Details

--- Time Elapsed: 349.591086ms ---



$ efa tenant po configure --name po1 --tenant tv3

PortChannel: po1 configured successfully.

--- Time Elapsed: 114.816703ms ---


(efa:extreme)extreme@node-1:~$ efa tenant po show
+------+--------+----+--------+-----+-------------+----------+---------+-----------------+------------+-------------
+-------------+
| Name | Tenant | ID | Speed  | MTU | Negotiation | Min Link |  Lacp   |      Ports      |   State
|  Dev State  |  App State  |
|      |        |    |        |     |             |  Count   | Timeout |                 |            |
|           |
+------+--------+----+--------+-----+-------------+----------+---------+-----------------+------------+-------------
+-------------+
| po1  |  tv3   | 1  | 10Gbps |     |   active    |    1     |  short  | 10.20.61.91[0/4] | po-created
```

```
| provisioned | cfg-in-sync |
|      |       |    |       |     |            |          |         | 10.20.61.90[0/4] |            |
|             |
+------+--------+----+-------+-----+------------+----------+---------+------------------+------------+-------------
+-------------+
Port Channel Details

--- Time Elapsed: 693.178441ms ---
```

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 table provides a list of port channel, device, and application state:

```
+---------------------------------+-------------+---------------+
|             State               |  Dev State  |   App State   |
|                                 |             |               |
+---------------------------------+-------------+---------------+
| po-min-links-count-set-pending | provisioned | cfg-refreshed |
|                                 |             |               |
+---------------------------------+-------------+---------------+
|       po-delete-pending         | provisioned | cfg-refreshed |
+---------------------------------+-------------+---------------+
|     po-port-delete-pending      | provisioned | cfg-refreshed |
+---------------------------------+-------------+---------------+
|           po-created            | provisioned |  cfg-in-sync  |
+---------------------------------+-------------+---------------+
|  po-lacp-timeout-set-pending    | provisioned | cfg-refreshed |
+---------------------------------+-------------+---------------+
|   po-description-set-pending    | provisioned | cfg-refreshed |
+---------------------------------+-------------+---------------+
|     po-mtu-delete-pending       | provisioned | cfg-refreshed |
+---------------------------------+-------------+---------------+
```

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.

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. When you create or update a port channel, you can provide an optional "min-link-count" for each XCO 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 the 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 prohibited 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 321
2. Update a Tenant VRF on page 326
3. Show a Tenant VRF on page 331
4. Delete a Tenant VRF on page 334
5. Shows a Tenant VRF Error on page 336
6. Configure a Tenant VRF on page 336
7. Distributed and Centralized Routing on page 361

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 |
128 } --max-path {1-128} | --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, 4.1.0*.

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 ---
```

Single Rack Deployment on SLX and OS ONE

- Single Rack Non-Clos Deployment (`single-rack-deployment = false`)
  - VRFs default to Distributed type with l3-extension enabled (cannot be disabled)
  - Centralized VRFs can be created with l3-extension disabled by default (can be enabled)
- Single Rack Non-Clos Deployment (`single-rack-deployment = true`)
  - All VRFs (Distributed and Centralized) have l3-extension disabled by default (cannot be enabled)
  - No "evpn irb ve <ve id> cluster-gateway", "bridge-domain add <bd id>", "map bridge-domain <bd id> vni <id>", "bridge domain <id>", and "interface Ve <id>" configurations on devices for VRFs with l3-extension disabled
  - Config Cleanup Behavior
    - When backup routing is enabled in a single rack deployment
      - Keeps backup routing neighbors, VE interface, and bridge-domain
      - Disables L3 extension configs, including (EVPN IRB under VRF, EVPN-BD mappings, and overlay-gateway mapping)
    - When backup routing is disabled in a single rack deployment
      - Removes L3 extensions, including EVPN IRB under VRF, EVPN-BD mapping, overlay-gateway mapping, VE interface, and bridge-domain configurations
- In a single-rack fabric, L3 extension is automatically disabled. As a result, configurations such as Suppress-ARP and ND - which were previously applied to L3 extension-enabled VRFs - will not be provisioned.

Upgrade Handling for Single-Rack Fabric on SLX

You can set single-rack-deployment value to true in two ways:

1. Before XCO upgrade: Setting it before XCO upgrade removes inapplicable layer3-extension configs (EVPN IRB BD, VE, etc.) during upgrade.
2. After upgrade: Setting it post upgrade (upgrade first, then set it to true) removes inapplicable layer3-extension configs (EVPN IRB BD, VE, etc.) post-upgrade.

Update a Tenant VRF

You can update an existing VRF for a tenant.

About This Task
Follow this procedure to update a tenant VRF.

Note

- For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.
- To update tenant VRF configuration, use the **`efa tenant vrf update`** command with a single `--operation` flag at a time. Multiple flags may cause ignored configurations.
- For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

You can update a VRF before or after you create an EPG. But, the VRF will reflect on the switches only after you create an EPG.

You can update the following operations:

- 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
- graceful-restart-update

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 | 128 } --max-path
```

```
{1-128} | -- 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, 4.1.0*.

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 table shows the application and device state of all VRFs:

```
+-------------------------------------------+--------------------+-------------+
```

```
|                    State                      |      Dev State      |  App State  |
|                                               |                     |             |
+-----------------------------------------------+---------------------+-------------+
|              vrf-device-created               |   not-provisioned   |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
|                 vrf-created                   |   not-provisioned   |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
|    vrf-device-static-route-delete-pending     |   not-provisioned   |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
|  vrf-device-static-route-bfd-delete-pending   |   not-provisioned   |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
|   vrf-device-network-route-delete-pending     |   not-provisioned   |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
| vrf-device-aggregate-address-delete-pending   |   not-provisioned   |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
| vrf-device-static-network-bfd-delete-pending | provisioning-failed |  cfg-ready  |
+-----------------------------------------------+---------------------+-------------+
|              vrf-device-created               |     provisioned     | cfg-in-sync |
+-----------------------------------------------+---------------------+-------------+
```

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 Port Channel, VRF, or EPG delete operations, use a single-name format for all entities within a tenant.
- For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*

.

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 ---
```

Delete Pending VRF Configuration

You can delete pending configuration on a VRF.

About This Task

Follow this procedure to push or remove the pending configuration on a VRF.

Procedure

Run the following command:

```
efa tenant vrf configure
```

The **`efa tenant vrf configure`** command pushes or removes the pending configuration of a VRF when it is in one of the following states:

`vrf-device-static-route-delete-pending | vrf-device-static-route-bfd-delete-pending | vrf-device-network-route-delete-pending | vrf-device-aggregate-address-delete-pending | vrf-device-local-asn-delete-pending | vrf-device-max-path-delete-pending | vrf-device-redist-delete-pending | vrf-device-rh-max-path-delete-pending | vrf-device-static-network-delete-pending`

Example

```
efa tenant vrf show --tenant t1
+------+--------+-------------+-------------+------------+--------------+------+-------
+--------+--------------------------------------+-----------------+-----------+
| Name | Tenant | Routing     | Centralized | Enable L3  | Redistribute | Max
| Local | Enable |               State               |     Dev State   | App State |
|      |        | Type        | Routers     | Extension  |              | Path
| Asn   | GR     |                                    |                |           |
+------+--------+-------------+-------------+------------+--------------+------+-------
+--------+--------------------------------------+-----------------+-----------+
|  v1  |   t1   | distributed |             |    true    |  connected   |
8   |       | false  | vrf-device-local-asn-delete-pending | not-provisioned | cfg-ready |
+------+--------+-------------+-------------+------------+--------------+------+-------
+--------+--------------------------------------+-----------------+-----------+

efa tenant vrf configure --tenant t1 --name v1
Vrf updated successfully.

--- Time Elapsed: 28.365307663s ---


efa tenant vrf show --tenant t1
+-------+--------+-------------+--------------+-------------+--------------+------+-------
+--------+--------------------+---------------+-------------+
| Name  | Tenant | Routing     | Centralized  |  Enable L3  | Redistribute
| Max   | Local | Enable |       State        |   Dev State   |  App State  |
|       |        | Type        | Routers      |  Extension  |              |
Path | Asn   | GR     |                    |               |             |
+-------+--------+-------------+--------------+-------------+--------------+------+-------
+--------+--------------------+---------------+-------------+
|  v1   |   t1   | distributed |              |     true    |  connected
|  8   |       | false  | vrf-device-created |  provisioned  | cfg-in-sync |
+-------+--------+-------------+--------------+-------------+--------------+------+-------
+--------+--------------------+---------------+-------------+
```

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, 4.1.0*.

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 using the **`efa tenant vrf configure [ --name | --tenant | --help ]`** command.

About This Task

Complete the following tasks to configure a tenant VRF in the XCO fabric:

Note

For syntax and command examples, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*

Procedure

1. Configure Local ASN on Tenant VRF on page 337
2. Enable Graceful Restart on Tenant VRF on page 343
3. Configure MaxPaths on Tenant VRF on page 344
4. Configure Resilient Hashing on Tenant VRF on page 345
5. Configure Redistribute Attribute on Tenant VRF on page 345

6. Configure Advertise Network and Static Network on Tenant VRF on page 347
7. Configure Aggregate Address on Tenant VRF on page 348
8. Configure EVPN IRB VE Cluster Gateway on a Tenant VRF on page 350
9. Route Distinguisher (RD) Allocation Independent of Route Target (RT) on page 355
10. Configure Static VRF Route on page 356
11. Configure BFD on Static VRF Route on page 357
12. Configure Backup Routing on Tenant VRF on page 358
13. Distributed and Centralized Routing on page 361
14. BFD Timers for Router BGP BFD and Static Route BFD Sessions on page 370
15. Configure Next Hop Recursion on page 371
16. Configure ECMP Paths on page 374
17. Enable Default Information Originate on page 379

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 and 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 |     | s1x-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

- Maximum value of `max-path` is 128.
- 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
```

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 | 128>
   ```

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 | 128>
   ```

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, 4.1.0*.

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
    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
  !
!
```

```
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::20/126 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. If the VRF subnets are distributed across multiple fabrics, you must enable the Layer3-extension.

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
!
```

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 370.

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-Extreme 8720-173# show running-config bridge-domain 3001 bridge-domain 3001 p2mp pw-profile default router-interface Ve 7001 bpdu-drop-enable ! leaf-Extreme 8720-173# 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-Extreme 8720-173# 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-Extreme 8720-173# 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-Extreme 8720-173# show running-config bridge-domain 3002 bridge-domain 3002 p2mp pw-profile default router-interface Ve 7002 bpdu-drop-enable ! leaf-Extreme 8720-173# 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-Extreme 8720-173# 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-Extreme 8720-173# 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.

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.

Note

The following rules apply to the Centralized Router (CR) configuration:

- A BL pair configured as CR excludes other BL pairs or single-homed BLs from being configured as CR.
- A single-homed BL configured as CR excludes other single-homed BLs or MCT pairs from being configured as CR.

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.

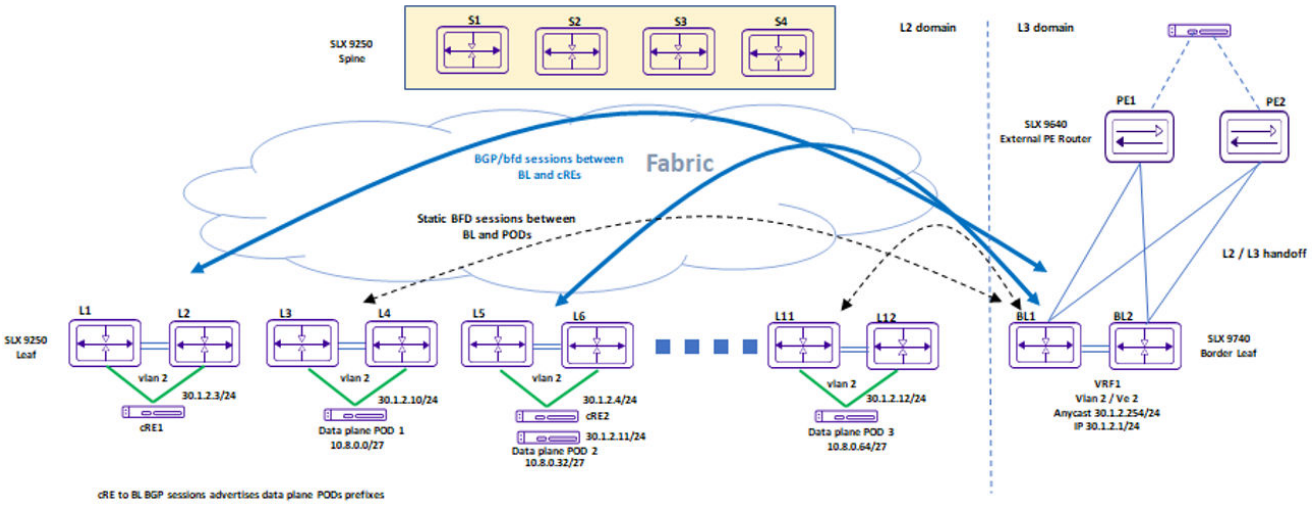

When a centralized routing is enabled for a given tenant VRF:

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>
      --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true/false>
```

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
    --pg-ipv6-uc-nbr-activate 10.20.246.29,v1:true
```

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>
    --pg-ipv6-uc-nbr-activate 10.20.246.29,v1:true
```

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
     --pg-ipv6-uc-nbr-activate 10.20.246.29,v1:true
```

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-deploymen 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".
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#
```

Configure ECMP Paths

You can configure ECMP and RH ECMP paths.

About This Task

Follow this procedure to configure ECMP and RH ECMP maximum paths to 128.

Procedure

To configure the EMP and RH ECMP maximum paths to 128, run the following command:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name>
                --max-path <1-128> --rh-max-path <8|16|64|128>

efa tenant vrf update -name <vrf-name> --tenant <tenant-name>
 --operation <max-path-add | max-path-delete | rh-max-path-add | rh-max-path-add>
 --max-path <1-128> --rh-max-path <8|16|64|128>
```

Example

```
efa:root)root@admin01:~# efa fabric show --name fabric1

Fabric Name: fabric1, Fabric Description: , Fabric Type: non-clos, Fabric Status: configure-success,
Fabric Health: Green
+-------------+-------+-----------+------------+------+--------------+-------------+-------------------
+-----------------+---------+-------+
| IP ADDRESS  | RACK  | HOST NAME |    ASN     | ROLE | DEVICE STATE |  APP STATE  | CONFIG GEN REASON
| PENDING CONFIGS | VTLB ID | LB ID |
+-------------+-------+-----------+------------+------+--------------+-------------+-------------------
+-----------------+---------+-------+
| 10.20.246.1 | rack1 | NH-1      | 4200000000 | Leaf | provisioned  | cfg in-sync | NA
| NA              | 2       | 1     |
| 10.20.246.2 | rack1 | NH-2      | 4200000000 | Leaf | provisioned  | cfg in-sync | NA
| NA              | 2       | 1     |
+-------------+-------+-----------+------------+------+--------------+-------------+-------------------
+-----------------+---------+-------+

(efa:root)root@admin01:~# efa tenant show
+------+---------+------------+-------------+-------------+-----------+-----------
+---------------------+--------------------------+
| Name |  Type   | VLAN Range | L2VNI Range | L3VNI Range | VRF Count | Enable BD |
Ports        | Mirror Destination Ports |
+------+---------+------------+-------------+-------------+-----------+-----------
+---------------------+--------------------------+
| ten1 | private |   11-20    |             |             |    10     |   false   |
10.20.246.1[0/1-10] |                          |
|      |         |            |             |             |           |           |
10.20.246.2[0/1-10] |                          |
+------+---------+------------+-------------+-------------+-----------+-----------
+---------------------+--------------------------+

(efa:root)root@admin01:~# efa tenant vrf create --name ten1vrf1 --tenant ten1

(efa:root)root@admin01:~# efa tenant vrf create --name ten1vrf2 --tenant ten1 --max-path 128 --rh-max-
path 128

(efa:root)root@admin01:~# efa tenant vrf show --detail
================================================================================================================
=====================================================
Name                    : ten1vrf1
Tenant                  : ten1
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: false
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

================================================================================================================
================================================
================================================================================================================
================================================
Name                     : ten1vrf2
Tenant                   : ten1
Routing Type             : distributed
Centralized Routers      :
Enable Layer3 Extension  : true
Redistribute             : connected
Max Path                 : 128
Local Asn                :
L3VNI                    :
EVPN IRB BD              :
EVPN IRB VE              :
BR VNI                   :
BR BD                    :
BR VE                    :
RH Max Path              : 128
Enable RH ECMP           : false
Enable Graceful Restart  : false
Enable NextHop Recursion: false
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:root)root@admin01:~# efa tenant epg create --name ten1epg1 --tenant ten1 --port 10.20.246.1[0/1]
--switchport-mode trunk --ctag-range 11 --vrf ten1vrf1 --anycast-ip 11:10.1.1.11/24
(efa:root)root@admin01:~# efa tenant epg create --name ten1epg2 --tenant ten1 --port 10.20.246.1[0/2]
--switchport-mode trunk --ctag-range 12 --vrf ten1vrf2 --anycast-ip 12:10.1.1.12/24


(efa:root)root@admin01:~# efa tenant vrf show --detail
================================================================================================================
================================================
Name                     : ten1vrf1
Tenant                   : ten1
Routing Type             : distributed
```

```
Centralized Routers     :
Enable Layer3 Extension : true
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
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

================================================================================================================
=====================================================
Name                    : ten1vrf2
Tenant                  : ten1
Routing Type            : distributed
Centralized Routers     :
Enable Layer3 Extension : true
Redistribute            : connected
Max Path                : 128
Local Asn               :
L3VNI                   : 8191
EVPN IRB BD             : 4095
EVPN IRB VE             : 8191
BR VNI                  : 4096
BR BD                   :
BR VE                   :
RH Max Path             : 128
Enable RH ECMP          : false
Enable Graceful Restart : false
Enable NextHop Recursion: false
Route Target            : import 102:102
                        : export 102:102
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

================================================================================================================
=====================================================

efa tenant vrf update --name ten1vrf1 --tenant ten1 --operation max-path-add --max-path 128
efa tenant vrf update --name ten1vrf1 --tenant ten1 --operation rh-max-path-add --rh-max-path 128
```

```
(efa:root)root@admin01:~# efa tenant vrf show --detail
================================================================================================================================
=======================================================
Name                    : ten1vrf1
Tenant                  : ten1
Routing Type            : distributed
Centralized Routers     :
Enable Layer3 Extension : true
Redistribute            : connected
Max Path                : 128
Local Asn               :
L3VNI                   : 8192
EVPN IRB BD             : 4096
EVPN IRB VE             : 8192
BR VNI                  : 4096
BR BD                   :
BR VE                   :
RH Max Path             : 128
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

================================================================================================================================
=======================================================
Name                    : ten1vrf2
Tenant                  : ten1
Routing Type            : distributed
Centralized Routers     :
Enable Layer3 Extension : true
Redistribute            : connected
Max Path                : 128
Local Asn               :
L3VNI                   : 8191
EVPN IRB BD             : 4095
EVPN IRB VE             : 8191
BR VNI                  : 4096
BR BD                   :
BR VE                   :
RH Max Path             : 128
Enable RH ECMP          : false
Enable Graceful Restart : false
Enable NextHop Recursion: false
Route Target            : import 102:102
                        : export 102:102
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
```

```
================================================================================================================
==================================================
```

Enable Default Information Originate

You can enable DIO (Default Information Originate) when you create or update a VRF. The "Default Information Originate" feature is commonly used in interior dynamic routing protocols such as iBGP. It is an important feature in network design as it ensures connectivity to networks that may not be explicitly known to all routers in the network, like the Internet. The main purpose of this feature is to allow a router to advertise a default route to other routers in the network, which is a 'catch-all' route that is used when a router does not have a more specific route for a destination in its routing table. This is represented as 0.0.0.0/0.

During Fabric creation (Clos and non-Clos), you can define the fabric wide setting to enable backup routing. When enabled, XCO creates an iBGP session between MCT Cluster peer to announce routes for traffic forwarding through the MCT peer in case all links to uplink from a given switch is lost. The creation of iBGP session is automated as part of the XCO tenant service. iBGP session announces routes correctly based on EBGP non-default static routes. If you rely on static routing, especially, default static route for routing towards the service provider router, the "default static router" is typically not announced to the MCT peer via iBGP session. This may result in traffic loss when uplinks are lost of specific MCT member.

To solve this issue, "Enable Default Information Originate" allows MCT peers to announce default routes explicitly. This can be done by configuring BGP peering with SLX configuration option "default-information-originate" at per VRF level.

About This Task

Follow this procedure to enable or disable DIO on a tenant VRF when you create or update a VRF.

- When you create or update a tenant VRF, you can choose the option to enable DIO. If you choose to enable DIO, XCO will configure DIO on the switches.
- 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.
- When you trigger VRF update operation, VRF is updated on the switches based on the endpoints present in the EPGs.
- When you configure a VRF, the DIO automatically gets configured.

Note

- By default, the DIO is disabled.
- When you upgrade from pre-XCO 3.4.0 to XCO 3.4.0 or later, DIO is disabled on VRFs and you can enable it.
- For information on hardware support, refer to the SLXOS documentation.

Procedure

1. Run the following command to enable or disable DIO when you create a VRF:

```
efa tenant vrf create --name <vrf-name> --tenant <tenant-name> --default-information-originate-enable {true|false}
```

2. Run the following command to enable or disable DIO when you update a VRF:

```
efa tenant vrf update --name <vrf-name> --tenant <tenant-name>
                --operation default-information-originate-update --default-information-
originate-enable {true|false}
```

Example

```
efa tenant vrf create --name vs --tenant t1 -- default-information-originate-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       : false
Default Information Originate  : 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       : false
Default Information Originate : 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
  default-information-originate
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  default-information-originate
  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
  default-information-originate
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  default-information-originate
  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 --
operation default-information-originate-update


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
Default Information Originate : 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
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  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
  redistribute connected
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf vs
  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.

Note

Using multiple `--ctag-range` options in a single **`efa tenant epg create`** or **`efa tenant epg update`** command causes an "anycast-ip not in ctag-range" error. To avoid this, specify a single contiguous range such as `--ctag-range <start-ctag>-<end-ctag>` format (for example, `--ctag-range 1101-1109`) to specify multiple tags instead of repeating `--ctag-range` with individual values.

Procedure

1. Create a Tenant Endpoint Group on page 383
2. Update a Tenant Endpoint Group on page 384
3. Show a Tenant Endpoint Group on page 387
4. Delete a Tenant Endpoint Group on page 387
5. Configure Network Property Description on Tenant EPG on page 397
6. Enable or Disable ICMP Redirect on Tenant EPG Networks on page 425
7. Update Anycast IP on an Existing Tenant Network on page 441
8. Configure Multiple Anycast IP on page 443
9. Configure IPv6 Neighbor Discovery (ND) on a Tenant Network on page 446
10. Configure BFD Session Type for an Endpoint Group on page 449
11. Configure Cluster Edge Port (CEP) Cluster Tracking for Endpoint Groups on page 449
12. Configure Suppress Address Resolution Protocol and Neighbor Discovery on VLAN or Bridge Domain on page 450
13. Configure Local IP for Endpoint Group on page 454
14. Software BFD Session Support on CEP on page 459
15. Enable Bulk Support for Tenant EPG APIs on page 468

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, 4.1.0*.

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.

Note

For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

For syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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 shows add and delete operation of a port group:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation port-group-add --po po2

EndpointGroup updated successfully.

--- Time Elapsed: 3.84922s ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation port-group-delete --po
po2

EndpointGroup updated successfully.

--- Time Elapsed: 1.1256584s ---
```

2. The following example shows add and delete operation of a port property:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation port-property-add --pp-
mac-acl-in ext-mac-permit-any-mirror-acl

EndpointGroup updated successfully.

--- Time Elapsed: 695.8661ms ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation port-property-delete
--pp-mac-acl-in ext-mac-permit-any-mirror-acl
```

```
EndpointGroup updated successfully.

--- Time Elapsed: 593.8088ms ---
```

3. The following example shows add and delete operation of a network property:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation network-property-add
--np-ip-acl-in 101:ext-ip-permit-any-mirror-acl

EndpointGroup updated successfully.

--- Time Elapsed: 3.014708s ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation network-property-
delete --np-ip-acl-in 101:ext-ip-permit-any-mirror-acl

EndpointGroup updated successfully.
```

4. The following example shows add and delete operation of a ctag range:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation ctag-range-add --ctag-
range 105 --anycast-ip 105:8.8.8.8/24

EndpointGroup updated successfully.

--- Time Elapsed: 3.7282495s ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation ctag-range-delete --
ctag-range 105 --anycast-ip 105:8.8.8.8/24

EndpointGroup updated successfully.

--- Time Elapsed: 1.4266126s ---
```

5. The following example shows add and delete operation of anycast IP:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation anycast-ip-add --
anycast-ip 101:3.3.3.3/24

EndpointGroup updated successfully.

--- Time Elapsed: 3.8720945s ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation anycast-ip-delete --
anycast-ip 101:3.3.3.3/24

EndpointGroup updated successfully.

--- Time Elapsed: 940.5274ms ---
```

6. The following example shows add and delete operation of a VRF:

   VRF deletion is independent of network policies and network properties.

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation vrf-delete

EndpointGroup updated successfully.

--- Time Elapsed: 185.5455ms ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation vrf-add --vrf v1 --
ctag-range 101 --anycast-ip 101:4.4.4.4/24

EndpointGroup updated successfully.
```

```
--- Time Elapsed: 6.9365858s ---
```

7. The following example shows add and delete operation of a local IP:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation local-ip-delete --
local-ip 101,10.20.246.3:1.10.1.1/24

EndpointGroup updated successfully.

--- Time Elapsed: 2.3330832s ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation local-ip-add --local-
ip 101,10.20.246.3:1.10.1.1/24

EndpointGroup updated successfully.

--- Time Elapsed: 840.6217ms ---
```

8. The following example shows add and delete operation of a DHCP relay address:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation dhcp-relay-address-ip-
add --dhcpv4-relay-address-ip 101,10.20.246.4:10.1.1.1

EndpointGroup updated successfully.

--- Time Elapsed: 2.552077s ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation dhcp-relay-address-ip-
delete --dhcpv4-relay-address-ip 101,10.20.246.4:10.1.1.1

EndpointGroup updated successfully.

--- Time Elapsed: 756.6756ms ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation dhcp-relay-address-ip-
add --dhcpv6-relay-address-ip 101,10.20.246.4:1::1

EndpointGroup updated successfully.

--- Time Elapsed: 760.959ms ---
```

9. The following example shows add and delete operation of a DHCP gateway:

```
$ efa tenant epg update --name e1 --tenant tenant11 --operation dhcp-relay-gateway-
ip-add --dhcpv6-relay-gateway-interface-ip 101,10.20.246.4:eth,0/5,3::3 --dhcpv6-relay-
gateway-interface 101,10.20.246.4:eth,0/5

EndpointGroup updated successfully.

--- Time Elapsed: 725.1882ms ---


$ efa tenant epg update --name e1 --tenant tenant11 --operation dhcp-relay-gateway-ip-
delete --dhcpv6-relay-gateway-interface-ip 101,10.20.246.4:eth,0/5,3::3 --dhcpv6-relay-
gateway-interface 101,10.20.246.4:eth,0/5

EndpointGroup updated successfully.

--- Time Elapsed: 768.0106ms ---
```

10. The following example creates a VLAN-based tenant with manual VNI mapping:

```
$ 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 ---
```

11. The following example creates a BD-based tenant:

```
$ 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 ---
```

12. The following example creates a tenant with auto-VNI with breakout ports:

```
$ 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 ---
```

13. The following creates a shared tenant:

```
$ 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, 4.1.0*.

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 the 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 delete a tenant endpoint group.

Note

- For Port Channel, VRF, or EPG delete operations, use a single-name format for all entities within a tenant.
- For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

To delete a tenant endpoint group, run the following command:

```
efa tenant epg delete --name ten1epg1 --tenant tenant1
```

Example

The following is an example output of deleting an EPG:

```
efa tenant epg show --detail
================================================================================================================================
=========================================================
Name          : e11
Tenant        : tv3
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.61.91[0/5]
              : 10.20.61.90[0/5]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : true
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 11
              : VRF                           : blue_dr
              : L3Vni                         : 14191

+-------------------------------+-------------------------------+-------------------------------
+-------------------------------+-------------------------------+
|          MAC ACL IN           |          MAC ACL OUT          |          IP ACL IN
|          IP ACL OUT           |          IPv6 ACL IN          |
+-------------------------------+-------------------------------+-------------------------------
+-------------------------------+-------------------------------+
| ext-mac-permit-any-mirror-acl | ext-mac-permit-any-mirror-acl | ext-ip-permit-any-mirror-acl | ext-
ip-permit-any-mirror-acl | ext-ipv6-permit-any-mirror-acl |
+-------------------------------+-------------------------------+-------------------------------
+-------------------------------+-------------------------------+
Port Property ACLs

+------------------+-------------+-------------+
|       Port       |  Dev State  |  App State  |
+------------------+-------------+-------------+
| 10.20.61.91[0/5] | provisioned | cfg-in-sync |
+------------------+-------------+-------------+
| 10.20.61.90[0/5] | 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 |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+-----------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
|  11  |   ctag-11   | 10002 |         | 11.11.11.9/29  | 11a::1a/125  |   T/T    | 10.20.61.90-
>11ab::31a/64, |      F/F      |  9000  |  9000   |      true      |     true     | provisioned | cfg-
in-sync |
|      |             |       |         |  11.11.11.1/29 | 11::11/125   |          |
1011::31a/64,        |               |        |         |                |              |
|             |
|      |             |       |         | 11.11.11.17/29 |              |          | 2.2.2.30/29,
2.2.2.132/29  |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 10.20.61.91-
>1011::32a/64, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 2.2.2.137/29,
2.2.2.37/29, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          |
11ab::32a/64        |               |        |         |                |              |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+-----------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------+---------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime |  Config Type  |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1002::/125   |    false     |    infinite    |      1020304       |   off-link    |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1003::/125   |    false     |    1020304     |      1020304       |   no-onlink   |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1004::/125   |    false     |    infinite    |      infinite      | no-autoconfig |
+------+----------------+--------------+----------------+--------------------+---------------+
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}         |                                         |
```

```
+------+-------------------------------------
+-------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------+
|  11  | 10.20.61.90->[{10.20.30.2,blue_dr}] |               10.20.61.90-
>[{11b::2,blue_dr}]                |         10.20.61.90->{10.20.30.1,,}         |         10.20.61.90-
>{eth,0/10,}          |
|      | 10.20.61.91->[{10.20.30.2,blue_dr}] |               10.20.61.91-
>[{11b::2,blue_dr}]                |      10.20.61.91->{10.20.30.1,eth,0/9}      |         10.20.61.91-
>{eth,0/10,}          |
+------+-------------------------------------
+-------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================================
=========================================================

--- Time Elapsed: 703.520561ms ---


efa tenant epg delete --name e11 --tenant tv3

EndpointGroup: e11 deleted successfully.

--- Time Elapsed: 170.028753ms ---
```

Delete Pending EPG Configuration

You can delete pending configuration on an EPG.

About This Task

Follow this procedure to push or remove the pending EPG configuration.

Procedure

Run the following command:

```
efa tenant epg configure
```

The **`efa tenant epg configure`** command pushes or removes the pending configuration for an EPG when it is in one of the following states:

```
epg-delete-pending | epg-port-group-delete-pending | epg-ctag-range-
delete-pending | epg-all-ctag-range-delete-pending | epg-vrf-delete-
pending | epg-dhcp-relay-address-ip-delete-pending | epg-dhcp-relay-
gateway-ip-delete-pending | epg-local-ip-delete-pending | epg-anycast-
ip-delete-pending | epg-port-property-delete-pending | epg-port-
property-update-pending | epg-network-property-delete-pending | epg-
network-property-update-pending state
```

Example

```
efa tenant epg show --detail
================================================================================================================
=========================================================
Name          : e11
```

```
Tenant        : tv3
Type          : extension
State         : epg-port-group-delete-pending
Description   :

Ports         : 10.20.61.91[0/5]
              : 10.20.61.90[0/5]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : true
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 11
              : VRF                           : blue_dr
              : L3Vni                         : 14191

+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
|          MAC ACL IN          |          MAC ACL OUT         |          IP ACL IN
|          IP ACL OUT          |          IPv6 ACL IN         |
+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
| ext-mac-permit-any-mirror-acl | ext-mac-permit-any-mirror-acl | ext-ip-permit-any-mirror-acl | ext-
ip-permit-any-mirror-acl | ext-ipv6-permit-any-mirror-acl |
+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
Port Property ACLs

+------------------+-------------+-------------+
|       Port       |  Dev State  |  App State  |
+------------------+-------------+-------------+
| 10.20.61.91[0/5] | provisioned | cfg-in-sync |
+------------------+-------------+-------------+
| 10.20.61.90[0/5] | 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 |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+------------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
|  11  |   ctag-11   | 10002 |         | 11.11.11.9/29  | 11a::1a/125  |   T/T    | 10.20.61.90-
>11ab::31a/64, |      F/F      |  9000  |  9000   |      true      |     true     | provisioned | cfg-
in-sync |
|      |             |       |         |  11.11.11.1/29 | 11::11/125   |          |
1011::31a/64,        |               |        |         |                |              |
|             |
|      |             |       |         | 11.11.11.17/29 |              |          | 2.2.2.30/29,
2.2.2.132/29  |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 10.20.61.91-
>1011::32a/64, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 2.2.2.137/29,
2.2.2.37/29, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          |
```

```
11ab::32a/64        |               |        |         |               |             |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+-----------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------+---------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime |  Config Type  |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1002::/125   |    false     |    infinite    |      1020304       |   off-link    |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1003::/125   |    false     |    1020304     |      1020304       |   no-onlink   |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1004::/125   |    false     |    infinite    |      infinite      | no-autoconfig |
+------+----------------+--------------+----------------+--------------------+---------------+
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}         |                                               |
+------+--------------------------------------
+-------------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------+
|  11  | 10.20.61.90->[{10.20.30.2,blue_dr}] |               10.20.61.90-
>[{11b::2,blue_dr}]                |         10.20.61.90->{10.20.30.1,,}         |         10.20.61.90-
>{eth,0/10,}          |
|      | 10.20.61.91->[{10.20.30.2,blue_dr}] |               10.20.61.91-
>[{11b::2,blue_dr}]                |      10.20.61.91->{10.20.30.1,eth,0/9}      |         10.20.61.91-
>{eth,0/10,}          |
+------+--------------------------------------
+-------------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================================
=========================================================

--- Time Elapsed: 703.520561ms ---


(efa:extreme)extreme@node-1:~$ efa tenant epg configure --name e11 --tenant tv3

EndpointGroup configured successfully.

--- Time Elapsed: 4.313882699s ---
(efa:extreme)extreme@node-1:~$
(efa:extreme)extreme@node-1:~$ efa tenant epg show --detail
```

```
================================================================================================================================
=========================================================
Name          : e11
Tenant        : tv3
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.61.91[0/5]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : true
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 11
              : VRF                           : blue_dr
              : L3Vni                         : 14191

+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
|          MAC ACL IN          |          MAC ACL OUT         |          IP ACL IN
|          IP ACL OUT          |          IPv6 ACL IN         |
+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
| ext-mac-permit-any-mirror-acl | ext-mac-permit-any-mirror-acl | ext-ip-permit-any-mirror-acl | ext-
ip-permit-any-mirror-acl | ext-ipv6-permit-any-mirror-acl |
+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
Port Property ACLs

+------------------+-------------+-------------+
|       Port       |  Dev State  |  App State  |
+------------------+-------------+-------------+
| 10.20.61.91[0/5] | 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 |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+------------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
|  11  |   ctag-11   | 10002 |         | 11.11.11.9/29  | 11a::1a/125  |   T/T    | 10.20.61.90-
>11ab::31a/64, |      F/F      |  9000  |  9000   |      true      |     true     | provisioned | cfg-
in-sync |
|      |             |       |         |  11.11.11.1/29 | 11::11/125   |          |
1011::31a/64,        |               |        |         |                |              |
|             |
|      |             |       |         | 11.11.11.17/29 |              |          | 2.2.2.30/29,
2.2.2.132/29  |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 10.20.61.91-
>1011::32a/64, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 2.2.2.137/29,
2.2.2.37/29, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          |
```

```
11ab::32a/64        |               |        |         |               |             |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+-----------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------+---------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime |  Config Type  |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1002::/125   |    false     |    infinite    |      1020304       |   off-link    |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1003::/125   |    false     |    1020304     |      1020304       |   no-onlink   |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1004::/125   |    false     |    infinite    |      infinite      | no-autoconfig |
+------+----------------+--------------+----------------+--------------------+---------------+
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}         |                                             |
+------+---------------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------+
|  11  | 10.20.61.90->[{10.20.30.2,blue_dr}] |               10.20.61.90-
>[{11b::2,blue_dr}]                |         10.20.61.90->{10.20.30.1,,}         |         10.20.61.90-
>{eth,0/10,}          |
|      | 10.20.61.91->[{10.20.30.2,blue_dr}] |               10.20.61.91-
>[{11b::2,blue_dr}]                |      10.20.61.91->{10.20.30.1,eth,0/9}      |         10.20.61.91-
>{eth,0/10,}          |
+------+---------------------------------------
+-----------------------------------------------------------------
+-----------------------------------------------+-----------------------------------------+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

=============================================================================================================
=========================================================

--- Time Elapsed: 753.003557ms ---
```

Force Delete an EPG

You can forcefully delete an EPG.

About This Task

Follow this procedure to forcefully delete an EPG.

Procedure

Run the following command:

```
efa tenant epg delete --name <> --teannt <> --force
```

The `force` option allows you to delete an EPG when dependent peer configuration is present. It deletes the EPG and associated dependent configuration from XCO and device.

The `force` option allows you to delete an EPG from the application when it is stuck in unexpected state and no other operations are allowed on it. Deleting an EPG with `force` option cleans up the entity from application and device.

Example

```
efa tenant epg show --detail
================================================================================================================================
=========================================================
Name           : e11
Tenant         : tv3
Type           : extension
State          : epg-port-group-delete-pending
Description    :

Ports          : 10.20.61.91[0/5]
               : 10.20.61.90[0/5]
POs            :
Port Property  : SwitchPort Mode               : trunk
               : Native Vlan Tagging           : true
               : Single-Homed BFD Session Type : auto
NW Policy      : Ctag Range                    : 11
               : VRF                           : blue_dr
               : L3Vni                         : 14191

+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
|          MAC ACL IN          |          MAC ACL OUT         |          IP ACL IN
|          IP ACL OUT          |          IPv6 ACL IN         |
+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
| ext-mac-permit-any-mirror-acl | ext-mac-permit-any-mirror-acl | ext-ip-permit-any-mirror-acl | ext-
ip-permit-any-mirror-acl | ext-ipv6-permit-any-mirror-acl |
+------------------------------+------------------------------+------------------------------
+------------------------------+------------------------------+
Port Property ACLs

+------------------+-------------+-------------+
|       Port       |  Dev State  |  App State  |
+------------------+-------------+-------------+
| 10.20.61.91[0/5] | provisioned | cfg-in-sync |
+------------------+-------------+-------------+
| 10.20.61.90[0/5] | 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 |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+-----------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
|  11  |   ctag-11   | 10002 |         | 11.11.11.9/29  | 11a::1a/125  |   T/T    | 10.20.61.90-
>11ab::31a/64, |      F/F      |  9000  |  9000   |      true      |     true     | provisioned | cfg-
in-sync |
|      |             |       |         |  11.11.11.1/29 | 11::11/125   |          |
1011::31a/64,        |               |        |         |                |              |
|             |
|      |             |       |         | 11.11.11.17/29 |              |          | 2.2.2.30/29,
2.2.2.132/29  |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 10.20.61.91-
>1011::32a/64, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          | 2.2.2.137/29,
2.2.2.37/29, |               |        |         |                |              |
|             |
|      |             |       |         |                |              |          |
11ab::32a/64        |               |        |         |                |              |
|             |
+------+-------------+-------+---------+----------------+--------------+----------
+-----------------------------+---------------+--------+---------+----------------+--------------
+-------------+-------------+
Network Property [Flags : * - Native Vlan]

+------+----------------+--------------+----------------+--------------------+---------------+
| Ctag | IPv6 ND Prefix | No Advertise | Valid Lifetime | Preferred Lifetime |  Config Type  |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1002::/125   |    false     |    infinite    |      1020304       |   off-link    |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1003::/125   |    false     |    1020304     |      1020304       |   no-onlink   |
+------+----------------+--------------+----------------+--------------------+---------------+
|  11  |   1004::/125   |    false     |    infinite    |      infinite      | no-autoconfig |
+------+----------------+--------------+----------------+--------------------+---------------+
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}         |                                        |
+------+--------------------------------------
+-------------------------------------------------------------
```

```
+-------------------------------------------+-------------------------------------------+
|  11  | 10.20.61.90->[{10.20.30.2,blue_dr}] |               10.20.61.90-
>[{11b::2,blue_dr}]                |         10.20.61.90->{10.20.30.1,,}         |         10.20.61.90-
>{eth,0/10,}          |
|      | 10.20.61.91->[{10.20.30.2,blue_dr}] |               10.20.61.91-
>[{11b::2,blue_dr}]                |      10.20.61.91->{10.20.30.1,eth,0/9}      |         10.20.61.91-
>{eth,0/10,}          |
+------+-------------------------------------
+-------------------------------------------------------------
+-------------------------------------------+-------------------------------------------+
DHCP Relay Ips


For 'unstable' entities, run 'efa tenant po/vrf show' for details

================================================================================================================
=====================================================

--- Time Elapsed: 703.520561ms ---



efa tenant epg delete --name e11 --tenant tv3 --force

EndpointGroup delete with "force" option will delete the device configuration corresponding to the
EndpointGroup and also deletes the EPG from the application. Do you want to proceed (Y/N): y

EndpointGroup: e11 deleted successfully.

--- Time Elapsed: 7.326435001s ---
```

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 does not 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.

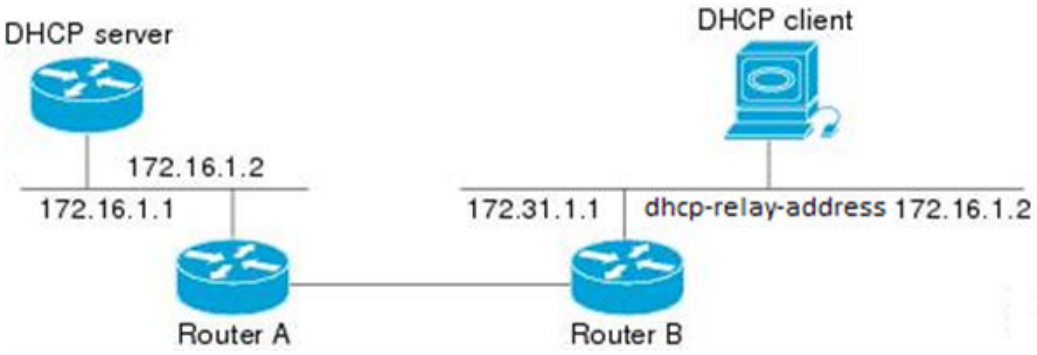

Figure 20: DHCP Client, DHCP Server, DHCP Relay

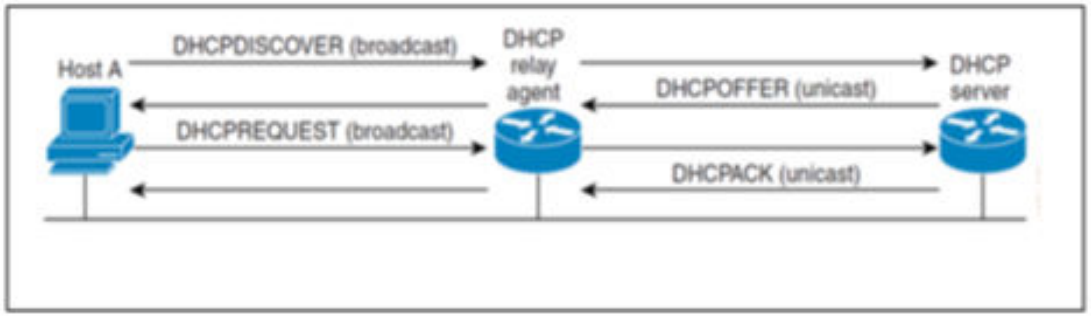

Figure 21: DHCP Relay Messages

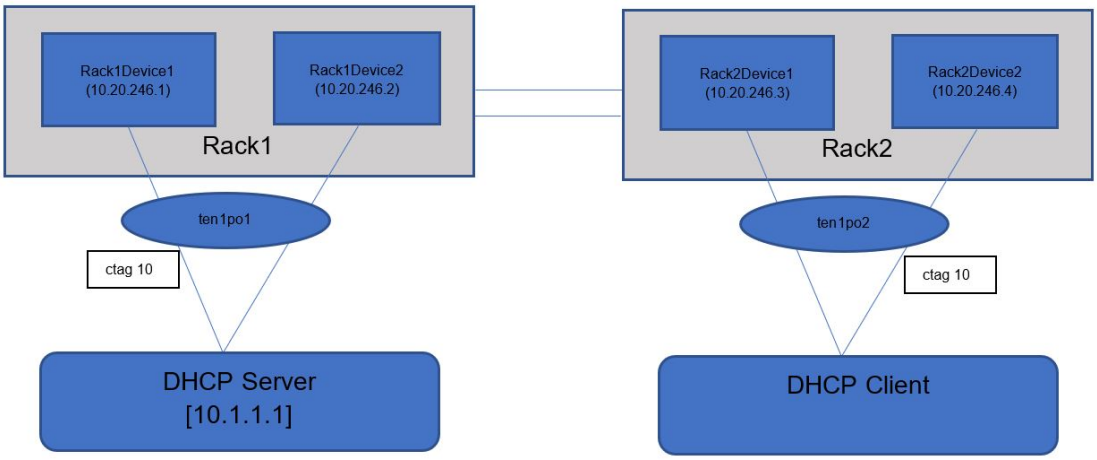

Figure 22: MultiRack Topology with DHCP Server, Client and Relay

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 Configuration	CLI Options	Description
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 Configuration	CLI Options	Description
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

Note

SLX requires the interface type and interface name for configuring the IPv6 DHCP relay gateway. XCO, in line with SLX, needs these two attributes to create or update operations.

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.

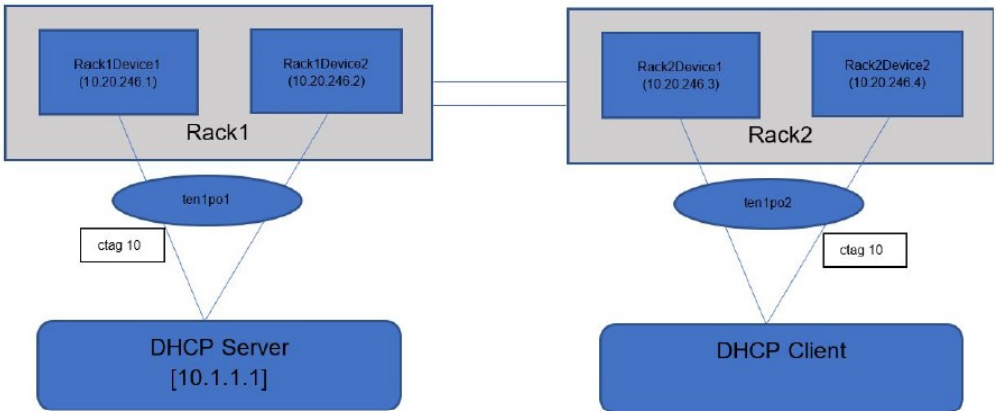

Figure 23: MultiRack Non-Clos Fabric with DHCP Server, Client, and Relay

```
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 Server 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 a DHCP Relay Gateway (`dhcp-relay-gateway-ip-add/delete`).

Use the following CLI 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**

Note
For information on enabling or disabling flooding for IP DHCP relay on a device, see Enable or Disable Flooding for IP DHCP Relay on page 817.

Enable or Disable ARP Suppression and Neighbor Discovery (ND)

You can enable or disable ARP Suppression and Neighbor Discovery (ND).

About This Task

Follow this procedure to enable or disable ARP Suppression and Neighbor Discovery (ND) on a tenant EPG when you configure network property during EPG update operations.

Note

- Latest value will be effective for EPG update and delete on suppress-arp and suppress-nd on common-ctag.
- You can only configure suppress-arp and suppress-ND attributes on a Layer 3 network EPG. It cannot be reconciled if an L2 VLAN or BD is configured OOB.

Procedure

Run the following command to enable or disable ARP Suppression and Neighbor Discovery (ND) on a tenant EPG:

When you delete a network property, a default value of "false" is configured.

```
efa tenant epg update --name <epg-name> --tenant <tenant-name>
         --operation  <network-propert-add | <network-property-update> | < network-
property-delete>
         --suppress-arp <ctag:suppress-arp> --suppress-nd <ctag:suppress-nd>
```

Example

- Layer 2 EPG Create

```
efa tenant epg create --name epg5 --tenant t1  --ctag-range 241 --switchport-mode
trunk-no-default-native --port 10.20.246.15[0/17],10.20.246.16[0/17]
EndpointGroup created successfully.
```

- Layer 2 EPG Update

```
efa tenant epg update --tenant t1 --name epg5 --operation network-property-update -
suppress-arp 241:true
Error:  Suppress ARP/ND configuration is not configured on L2 Network EPG
```

- Layer 3 EPG Create

```
efa tenant epg create --name epg1 --tenant t1 --switchport-mode trunk --port
10.20.246.15[0/12],10.20.246.16[0/12] --ctag-range 211-213 --vrf vrf11 --anycast-ip
211:33.1.1.1/24 --anycast-ipv6 212:500::10/31 --anycast-ipv6 213:600::10/31 --anycast-ip
213:43.1.1.2/24
```

```
efa tenant epg show --name epg1 --tenant t1 --detail
================================================================================================================
============================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/12]
              : 10.20.246.16[0/12]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 211-213
              : VRF                           : vrf11
              : L3Vni                         : 8160

+------------+-------------+-----------+------------+-------------+
| 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[0/12] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
| 10.20.246.16[0/12] | 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
|             |             |
+------+------------------------+-------+---------+--------------+--------------+----------
+----------------------+---------------+--------+---------+---------------+--------------
+-------------+-------------+
| 211  | Tenant L3 Extended VLAN |  211  |         | 33.1.1.1/24  |              |   T/F
|                      |      F/F      |        |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+------------------------+-------+---------+--------------+--------------+----------
+----------------------+---------------+--------+---------+---------------+--------------
+-------------+-------------+
| 212  | Tenant L3 Extended VLAN |  212  |         |              |  500::10/31  |   F/T
|                      |      F/F      |        |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+------------------------+-------+---------+--------------+--------------+----------
+----------------------+---------------+--------+---------+---------------+--------------
+-------------+-------------+
| 213  | Tenant L3 Extended VLAN |  213  |         | 43.1.1.2/24  |  600::10/31  |   T/T
|                      |      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
```

- Layer 3 EPG Update (network-property-add) with suppress-rap true and suppress-ND true

```
efa tenant epg update --tenant t1 --name epg1 --operation network-property-add -suppress-arp
211:true -suppress-nd 211:true
EndpointGroup updated successfully.


efa tenant epg show --name epg1 --tenant t1 --detail
=============================================================================================================
============================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/12]
              : 10.20.246.16[0/12]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 211-213
              : VRF                           : vrf11
              : L3Vni                         : 8160

+------------+-------------+-----------+------------+-------------+
| 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[0/12] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
| 10.20.246.16[0/12] | 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
|             |             |
+------+-------------------------+-------+---------+--------------+--------------+----------
+-----------------------+---------------+---------+---------+----------------+--------------
+-------------+-------------+
| 211  | Tenant L3 Extended VLAN |  211  |         | 33.1.1.1/24  |              |   F/T
|                       |      F/F     |        |          |     false      |    false     |
provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------+----------
+-----------------------+---------------+---------+---------+----------------+--------------
+-------------+-------------+
| 212  | Tenant L3 Extended VLAN |  212  |         |              |  500::10/31  |   F/T
|                       |      F/F      |        |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+-------------------------+-------+---------+--------------+--------------+----------
+-----------------------+---------------+---------+---------+----------------+--------------
+-------------+-------------+
| 213  | Tenant L3 Extended VLAN |  213  |         | 43.1.1.2/24  |  600::10/31  |   T/T
|                       |      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
```

```
================================================================================================
============================================================
```

- Layer 3 EPG Update (network-property-update) with suppress-rap true and suppress-ND false

```
efa tenant epg update --tenant t1 --name epg1 --operation network-property-update --suppress-arp
212:true -suppress-nd 213:false
EndpointGroup updated successfully.


efa tenant epg show --name epg1 --tenant t1 --detail
================================================================================================
============================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :
Ports         : 10.20.246.15[0/12]
              : 10.20.246.16[0/12]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 211-213
              : VRF                           : vrf11
              : L3Vni                         : 8160
+------------+-------------+-----------+------------+-------------+
| 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[0/12] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
| 10.20.246.16[0/12] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
Port Property States
+------+-------------------------+-------+------+-------------+------------+----------
+-----------------------+---------------+-----+---------+----------------+--------------
+-------------+-------------+
| Ctag |          Ctag           | L2Vni | BD   | Anycast     | Anycast    | Suppress |
Local IP        | Icmp Redirect | IP  | IPv6 ND |    IPv6 ND     |   IPv6 ND    |  Dev State  |
App State  |
|      |       Description       |       | Name | IPv4        | IPv6       |  ARP/ND  | [Device-IP-
>Local-IP] |   IPv4/IPv6   | MTU |   MTU   | Managed Config | Other Config |
|             |
+------+-------------------------+-------+------+-------------+------------+----------
+-----------------------+---------------+-----+---------+----------------+--------------
+-------------+-------------+
| 211  | Tenant L3 Extended VLAN |  211  |      | 33.1.1.1/24 |            |   F/T
|                        |      F/F      |     |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+-------------------------+-------+------+-------------+------------+----------
+-----------------------+---------------+-----+---------+----------------+--------------
+-------------+-------------+
| 212  | Tenant L3 Extended VLAN |  212  |      |             | 500::10/31 |   T/T
|                        |      F/F      |     |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+-------------------------+-------+------+-------------+------------+----------
+-----------------------+---------------+-----+---------+----------------+--------------
+-------------+-------------+
```

```
| 213  | Tenant L3 Extended VLAN |  213  |      | 43.1.1.2/24 | 600::10/31 |   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
```

- Layer 3 EPG Update (network-property-delete)

```
efa tenant epg update --tenant t1 --name epg1 --operation network-property-delete --suppress-arp
211:false -suppress-nd 211:false
EndpointGroup updated successfully.


efa tenant epg show --name epg1 --tenant t1 --detail
================================================================================================================
============================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         : 10.20.246.15[0/12]
              : 10.20.246.16[0/12]
POs           :
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 211-213
              : VRF                           : vrf11
              : L3Vni                         : 8160

+------------+-------------+-----------+------------+-------------+
| 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[0/12] | provisioned | cfg-in-sync |
+--------------------+-------------+-------------+
| 10.20.246.16[0/12] | 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
|             |             |
+------+------------------------+-------+---------+-------------+-------------+----------
+-----------------------+---------------+---------+---------+----------------+--------------
+-------------+-------------+
| 211  | Tenant L3 Extended VLAN |  211  |         | 33.1.1.1/24  |              |   F/F
|                       |      F/F      |        |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+------------------------+-------+---------+-------------+-------------+----------
+-----------------------+---------------+---------+---------+----------------+--------------
+-------------+-------------+
| 212  | Tenant L3 Extended VLAN |  212  |         |              |  500::10/31  |   T/T
|                       |      F/F      |        |         |     false      |    false     |
provisioned | cfg-in-sync |
+------+------------------------+-------+---------+-------------+-------------+----------
+-----------------------+---------------+---------+---------+----------------+--------------
+-------------+-------------+
| 213  | Tenant L3 Extended VLAN |  213  |         | 43.1.1.2/24  |  600::10/31  |   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
```

- SLX configuration

```
Rack1-Device1#   show   run   vlan
211-213
vlan 211
 router-interface Ve 211
  description  Tenant  L3  Extended
VLAN
!
vlan 212
 router-interface Ve 212
 suppress-nd
 suppress-arp
  description  Tenant  L3  Extended
VLAN
!
vlan 213
 router-interface Ve 213
 suppress-arp
  description  Tenant  L3  Extended
VLAN
!
```

```
Rack1-Device2#   show   run   vlan
211-213
vlan 211
 router-interface Ve 211
  description  Tenant  L3  Extended
VLAN
!
vlan 212
 router-interface Ve 212
 suppress-nd
 suppress-arp
  description  Tenant  L3  Extended
VLAN
!
vlan 213
 router-interface Ve 213
 suppress-arp
  description  Tenant  L3  Extended
VLAN
!
```

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 does not 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
 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
!
```

```
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
!
```

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) 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, 4.1.0*

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:	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 - 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:

```
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 - 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:	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 !	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:false

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 ---
```

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 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 !	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 !

Update Anycast IP on an Existing Tenant Network

You can add 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 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 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 the local IP when you create an EPG, run the following commands:

```
$ efa tenant epg create --tenant "bdST11" --name "bdST11epg1" --type l3-hand-off --
switchport-mode trunk-no-default-native --single-homed-bfd-session-type auto --port
10.64.224.19[0/5:2],10.64.224.18[0/5:2] --vrf bdST11vrf1 --ctag-range 201 --bridge-
```

```
domain 201:Auto-BD-4098 --ctag-description "201:Tenant L3 Hand-off BD" --local-ip
201,10.64.224.18:201.1.1.1/24 --local-ip 201,10.64.224.19:201.1.1.2/24


EndpointGroup created successfully.


--- Time Elapsed: 2.171307492s ---



$ efa tenant epg show --tenant "bdST11" --name "bdST11epg1" --detail

================================================================================================
================================================================================

Name           : bdST11epg1

Tenant         : bdST11

Type           : l3-hand-off

State          : epg-with-port-group-and-ctag-range

Description    :


Ports          : 10.64.224.18[0/5:2]

               : 10.64.224.19[0/5:2]

POs            :

Port Property : SwitchPort Mode               : trunk-no-default-native

               : Single-Homed BFD Session Type : auto

NW Policy      : Ctag Range                    : 201

               : VRF                           : bdST11vrf1


+------------+-------------+-----------+------------+-------------+

| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |

+------------+-------------+-----------+------------+-------------+

Port Property ACLs


+--------------------+-------------+-------------+

|        Port        |  Dev State  |  App State  |

+--------------------+-------------+-------------+

| 10.64.224.18[0/5:2] | provisioned | cfg-in-sync |

+--------------------+-------------+-------------+

| 10.64.224.19[0/5:2] | 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 |             |             |

+------+-----------------------+-------+--------------+--------------+--------------
+----------+-----------------------------+---------------+--------+---------
+----------------+--------------+-------------+-------------+

| 201  | Tenant L3 Hand-off BD | 4097  | Auto-BD-4098 |              |
|   F/F    | 10.64.224.18->201.1.1.1/24 |      F/F      |        |         |
false      |    false     | provisioned | cfg-in-sync |

|      |                       |       |              |              |
|          | 10.64.224.19->201.1.1.2/24 |               |        |
|               |              |             |             |

+------+-----------------------+-------+--------------+--------------+--------------
+----------+-----------------------------+---------------+--------+---------
+----------------+--------------+-------------+-------------+

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: 104.300999ms ---
```

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.

Software BFD Session Support on CEP

You can configure software Bidirectional Forwarding Detection (BFD) sessions on a Cluster Edge Port (CEP) on Extreme 8520 and Extreme 8720 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.

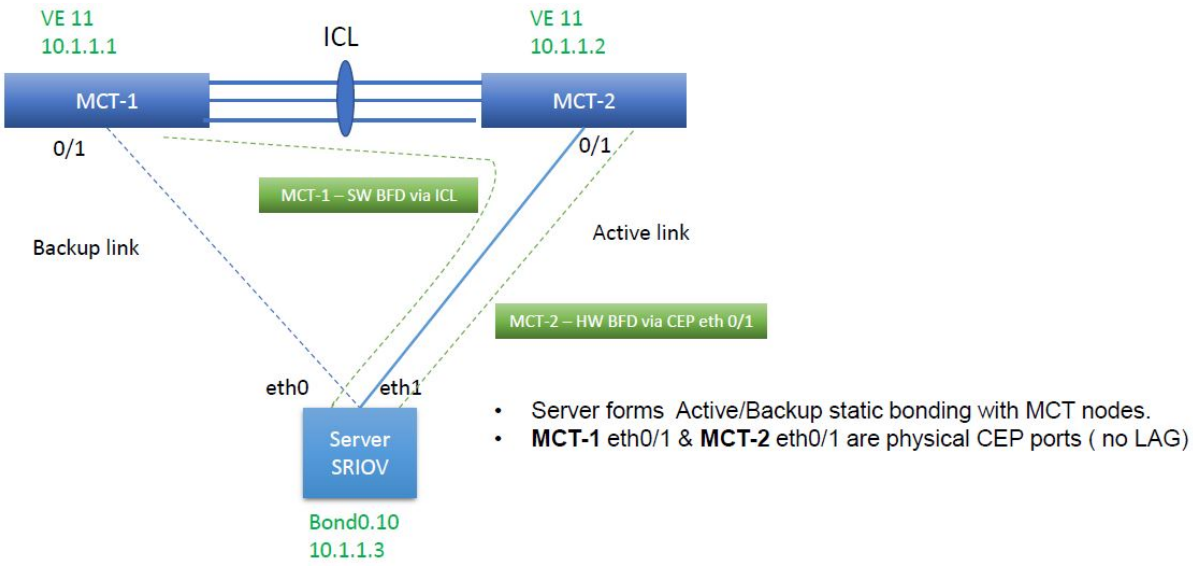

BFD Session Formation with SRIOV Server after Link Failover

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.

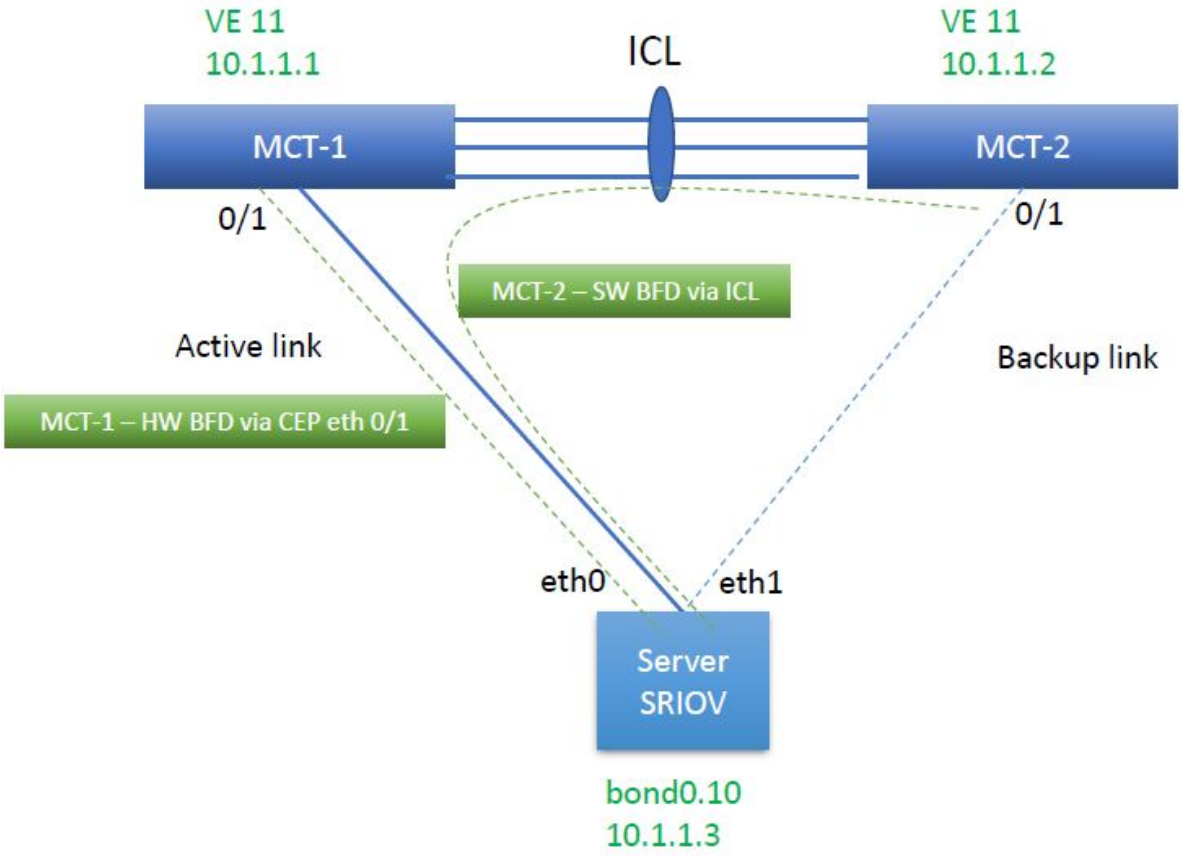

BFD Session Formation with SRIOV Server

- Limitation in Extreme 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 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 --single-homed-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 --single-homed-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 Extreme 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

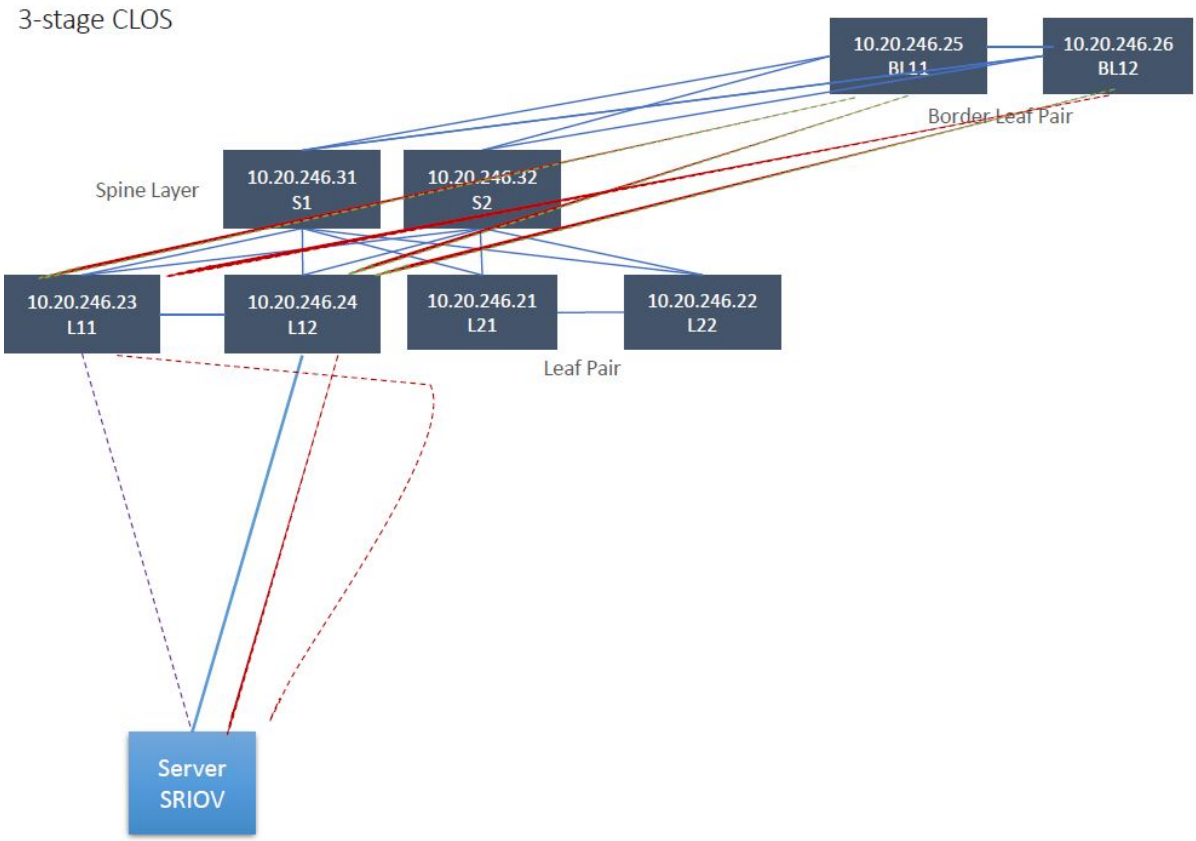

Bulk Support for Tenant EPG API

XCO 4.0.2 and later releases introduce bulk support for tenant EPG APIs, enabling faster configuration, updates, and deletions of EPG configurations by processing multiple EPGs in a single request.

The following are some important conditions for bulk support in tenant EPG APIs:

- Prior Release Limitation: In XCO versions prior to 4.0.2, bulk EPG port-group updates via a single API were only supported if all EPGs involved had Layer 2 configurations exclusively.
- CLI Support: CLI support for Bulk APIs is not available.
- Performance Improvement: Bulking enhances performance by consolidating individual API calls for all EPGs into a single request, significantly reducing overall processing time.
- Error Handling: If a failure occurs due to one EPG in the request, the entire request will fail with a 4xx or 5xx error.
- Layer 2 Support: Starting from XCO 3.8.0, bulk configuration is supported for Layer 2 type EPGs, allowing you to configure EPGs with CTAG ID and Port-Group in bulk.
- Layer 3 Support: Layer 3 type EPG configurations support bulking in XCO 4.0.2. The following configuration is supported:
  - EPG with VRF configuration and its corresponding attributes
  - EPG with Anycast IP Address configuration (IPv4 and IPv6)

  - EPG with Local IP Address configuration (IPv4 and IPv6)
  - EPG with no dhcp-relay, dhcp-server, and QoS configurations

  Layer 3 configuration for VRF, Anycast, and Local IP is optimized, and when combined with port group additions, these configurations are efficiently pushed to devices for bulked EPGs.
- EPG Create: It supports both bulk and non-bulk creation of EPGs. It supports bulk creation of up to 350 EPGs in a single request, and allows EPGs with either a single CTAG or a range of CTAGs. However, the bulk operation is not optimized for performance.
- EPG Update: Only the port-group-add and port-group-delete operations are supported for bulked EPG requests.
- EPG Delete: It supports both bulk and non-bulk deletion of EPGs and allows bulk deletion of up to 350 EPGs in a single request. However, the bulk operation is not optimized for performance.
- Global Configuration: After updating the Global configuration, restart tenant services to ensure proper functionality.

Limitations

Use this topic to learn about the limitations of bulk support for tenant EPG API.

Functionality

- Only VLAN-based Tenants are supported.
- Error handling is done at the API level.

  For bulk EPG updates in a single API, error responses are returned at the API level.
- Complete feature parity with non-bulk APIs is not available.

Idempotency

Idempotency is supported at the bulk EPG level, but it's applied individually to each EPG. In bulk requests, idempotency checks are performed for each EPG, and any EPGs that are idempotent will be skipped in subsequent operations.

Rollback on Failure

When handling bulk requests as a single transaction, any intermediate failure results in configuration failure. The entire bulk request is rolled back in case of any failure.

Maximum EPGs for Bulking Configuration

By default, the maximum number of EPGs supported in any bulk request is 350. You cannot configure this value either via CLI or REST API.

Enable Bulk Support for Tenant EPG APIs

About This Task

Maximum EPGs for Bulking: By default, the maximum number of EPGs supported in any bulk request is 350. You cannot configure this value either via CLI or REST API.

Follow the procedure to enable bulk support for tenant EPG APIs.

Procedure

1. To verify the default status of tenant EPG API bulking, run the following command:

   By default, the tenant EPG API bulking is disabled.

```
efa system feature show
+---------------------------------+----------+
|          FEATURE NAME           |  STATUS  |
+---------------------------------+----------+
| Inflight Transaction Auto       | Enabled  |
| Recovery                        |          |
+---------------------------------+----------+
| Tenant Api Concurrency          | Enabled  |
+---------------------------------+----------+
| Tenant EPG Api Bulking          | Disabled |
+---------------------------------+----------+
```

2. Enable or disable bulk support in tenant EPG APIs.

   You can enable or disable the bulk support using either the EFA CLI or the REST API. By default, bulk support is disabled in tenant EPG APIs.

   a. Using CLI: To enable the tenant API bulk operation, run the following command:

```
efa system feature update --bulk-epg-api Enable

Output:
Feature Setting Updated Successful
```

   b. Using API: To enable the tenant API bulk operation, use the following API:

```
curl -X PUT http://gosystem-service:8090/v1/system/feature  -H "Content-Type:
application/json" -d '{"keyval":[{"key":"BulkEPGAPI","value":"Enable"}]}'
```

   You can display the result by using the following API:

```
curl -X GET http://gosystem-service:8090/v1/system/feature

{"keyval":[{"key":"Inflight Transaction Auto Recovery","value":"Enabled"},
{"key":"Tenant Api Concurrency","value":"Enabled"},{"key":"Tenant EPG API
Bulking","value":"Enabled"}]}
```

3. Restart the tenant services.

   A restart of the tenant services is necessary for the changes to be implemented and take effect.

```
efactl restart-service gotenant-service
Are you sure you want to restart "gotenant-service"? [Y/n]
y
"gotenant-service" has been stopped
"gotenant-service" has been started
"gotenant-service" has been restarted
```

4. To verify that the tenant EPG API bulking is enabled, run the following command:

```
efa system feature show
+---------------------------------+---------+
|          FEATURE NAME           | STATUS  |
+---------------------------------+---------+
| Inflight Transaction Auto       | Enabled |
| Recovery                        |         |
+---------------------------------+---------+
| Tenant Api Concurrency          | Enabled |
+---------------------------------+---------+
| Tenant EPG Api Bulking          | Enabled |
+---------------------------------+---------+
```

```
--- Time Elapsed: 228.076737ms ---
(efa:extreme)extreme@tpvm101:~$
```

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 471
2. Create BGP Dynamic Peer on page 479
3. Getting the Operational State of the BGP Peers on page 488
4. Configure Route Map Attribute on page 490
5. Configure remove-private-as on BGP Peer on page 494
6. Configure default-originate to Advertise Default Route on BGP Peer on page 497
7. Configure Backup Routing Neighbors on BGP Peer on page 501
8. Configure Send-Community on Tenant BGP Peer on page 502
9. Configure Out-of-band for a Tenant BGP Peer or Peer Group on page 505
10. Configure Multi Protocol BGP on Tenant Dynamic BGP Peer on page 516
11. Configure Multi Protocol BGP on Tenant Static BGP Peer on page 508
12. Enable or Disable MP BGP Capability for IPv6 Prefix Exchange over IPv4 Peer on page 517

Note

Consuming a BGP peer requires prior consumption of the peer group via 'Service BGP', which is the expected design in the Extreme OS ONE.

To consume BGP peers on Extreme OS ONE devices, follow these steps

The procedure applies to both OOB peers and backup routing peers.

a. Create an L3 EPG with backup routing enabled to push the BR PG and peer to the device.
b. Consume the BR PG using the **`Service BGP PG`** command.
c. Consume the BR peer using the **`Service BGP Peer`** command.

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.

Note

For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

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, 4.1.0*

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.

Note

For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

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
 !
!
```

Delete Pending BGP Peer Configuration

You can delete pending configuration on a BGP peer.

About This Task

Follow this procedure to remove the pending configuration on a BGP peer.

Procedure

Run the following command:

```
efa tenant service bgp peer configure
```

The **`efa tenant service bgp peer configure`** command pushes or removes a pending configuration for a BGP peer instance when it is in one of the following states:

`bgp-peer-delete-pending | bgp-peer-peer-delete-pending`

Example

```
efa tenant service bgp peer show
================================================================================================
================================================================================================
========
Name        : customer_2
Tenant      : tv3
State       : bgp-peer-peer-delete-pending

+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
|  Device IP  |   VRF   | AFI  |  SAFI   | Remote IP | Remote ASN | Activate | Next Hop
|   Update    |   BFD   |           BFD            |  Dev State  |  App State  |
|             |         |      |         |           |            |          |   Self
|  Source IP  | Enabled | [Interval,Rx,Multiplier] |             |             |
+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
| 10.20.61.91 | blue_dr | ipv4 | unicast | 1.1.1.10  |   95001    |   true   |  always  |
10.11.12.13 |  true   |       50, 5000, 50       | provisioned | cfg-in-sync |
+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
| 10.20.61.90 | blue_dr | ipv4 | unicast | 1.1.1.10  |   95001    |   true   |  always  |
10.11.12.13 |  true   |       50, 5000, 50       | provisioned | cfg-in-sync |
+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
Static Peer Details

[Flag '*' indicates Multi protocol capability]


+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
|  Device-IP  |   VRF   | AFI  |  SAFI   | Listen Range  | Listen | Peer Group |  Dev
State  |  App State  |
|             |         |      |         |               | Limit  |
|             |             |
+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
| 10.20.61.91 | blue_dr | ipv4 | unicast | 15.16.16.0/28 |   10   |    pg1     |
provisioned | cfg-in-sync |
+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
| 10.20.61.90 | blue_dr | ipv4 | unicast | 15.16.16.0/28 |   10   |    pg1     |
```

```
provisioned | cfg-in-sync |
+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
Dynamic Peer Details

================================================================================================
================================================================================================
========

--- Time Elapsed: 341.331949ms ---

(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer configure --name customer_2 --
tenant tv3

BgpService configured successfully.

--- Time Elapsed: 3.734570672s ---
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer show
================================================================================================
================================================================================================
========
Name        : customer_2
Tenant      : tv3
State       : bgp-peer-created

+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
|  Device IP  |   VRF   | AFI  |  SAFI   | Remote IP | Remote ASN | Activate | Next Hop
|   Update    |   BFD   |           BFD            |  Dev State  |  App State  |
|             |         |      |         |           |            |          |   Self
|  Source IP  | Enabled | [Interval,Rx,Multiplier] |             |             |
+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
| 10.20.61.91 | blue_dr | ipv4 | unicast | 1.1.1.10  |   95001    |   true   |  always  |
10.11.12.13 |  true   |       50, 5000, 50       | provisioned | cfg-in-sync |
+-------------+---------+------+---------+-----------+------------+----------+----------
+-------------+---------+--------------------------+-------------+-------------+
Static Peer Details

[Flag '*' indicates Multi protocol capability]


+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
|  Device-IP  |   VRF   | AFI  |  SAFI   | Listen Range  | Listen | Peer Group |  Dev
State  |  App State  |
|             |         |      |         |               | Limit  |
|             |             |
+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
| 10.20.61.91 | blue_dr | ipv4 | unicast | 15.16.16.0/28 |   10   |    pg1     |
provisioned | cfg-in-sync |
+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
| 10.20.61.90 | blue_dr | ipv4 | unicast | 15.16.16.0/28 |   10   |    pg1     |
provisioned | cfg-in-sync |
+-------------+---------+------+---------+---------------+--------+------------
+-------------+-------------+
Dynamic Peer Details


================================================================================================
================================================================================================
========
```

```
--- Time Elapsed: 425.060566ms ---
```

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 <username> --
password <password>
efa inventory device register --ip 10.20.246.21,10.20.246.22 --username <username> --
password <password>
efa inventory device register --ip 10.20.246.14 --username <username> --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
<username> --password <password>
efa fabric device add --ip 10.20.246.23 --role leaf  --name fabric4 --username
<username> --password <password>
efa fabric device add --ip 10.20.246.24 --role leaf  --name fabric4 --username
<username> --password <password>
efa fabric device add --ip 10.20.246.22 --role border-leaf --name fabric4 --username
<username> --password <password>
efa fabric device add --ip 10.20.246.21 --role border-leaf --name fabric4 --username
<username> --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
   --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true>

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
   --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true
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, 4.1.0*.

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, 4.1.0*.

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 :	: v1	:
Dev	AFI	Remove Private
State	: ipv4	AS : false
: provisioned	SAFI	Default
App	: unicast	Originate :
State	Remote	true
: cfg-in-sync	IP :	Default Originate
	10.20.30.40	Route Map : tt
	Remote	Prefix List
	ASN :	In :
	50000	Prefix List
	Next Hop	Out :
	Self :	Route Map
	false	In :
	Update Source	Route Map
	IP :	Out :
	BFD	Dev
	Enabled	State
	: true	: provisioned
	BFD	App
	Interval	State
	: 0	: cfg-in-sync
	BFD	
	Rx	
	: 0	

```
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, 4.1.0*.

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
    --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true/false>

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
    --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true/false>

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
    --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true/false>

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

================================================================================
========
```

Multi Protocol BGP

You can configure multi protocol BGP on BGP static and dynamic peers.

XCO supports IPv4 BGP session which can carry IPv6 prefixes.

This allows you to scale IPv4 or IPv6 dual stack without creating an IPv6 or IPv4 BGP or BFD session. This is applicable for both **static** and **dynamic** BGP peers.

Dynamic BGP peers are always associated with a **BGP peer-group**.

Configure Multi Protocol BGP on Tenant Static BGP Peer

You can configure multi protocol BGP.

About This Task
Follow this procedure to configure multi protocol BGP on tenant static BGP peer.

The session specific configurations of an IPv4 BGP neighbor are placed under the IPv4 address-family. The IPv4 route (prefix) specific configurations of an IPv4 BGP neighbor are placed under the IPv4 address-family. The IPv6 route (prefix) specific configurations of an IPv4 BGP neighbor are placed under the IPv6 address-family.

You can enable multi protocol BGP on a static peer by activating an IPv4 neighbor under the IPv6 address-family.

Note
Backup routing neighbors do not support multi protocol BGP.

Procedure

1. On SLX devices, run the following commands under IPv4 or IPv6 address-family:
   a. IPv4 Address-family
      - Session Specific Configuration

```
NH-Leaf1(config-bgp-ipv4u-vrf)# neighbor 25.x.x.x ?
Possible completions:

  advertisement-interval   Minimum interval between sending BGP routing updates
  announce-rpki-state      Announce RPKI state
  as-override              Override matching AS-number while sending update
  bfd                      Enable BFD session for the neighbor
  description              Neighbor by description
  ebgp-btsh                Enable EBGP TTL Security Hack Protection
  ebgp-multihop            Allow EBGP neighbors not on directly connected
networks
  enforce-first-as         Enforce the first AS for EBGP routes
  graceful-restart         Enable graceful restart for the neighbor
  local-as                 Assign local-as number to neighbor
  maxas-limit              Impose limit on number of ASes in AS-PATH attribute
  next-hop-self            Disable the next hop calculation for this neighbor
  password                 Enable TCP-MD5 password protection
  remote-as                Specify a BGP neighbor
  remove-private-as        Remove private AS number from outbound updates
  shutdown                 Administratively shut down this neighbor
  soft-reconfiguration     Per neighbor soft reconfiguration
  static-network-edge      Neighbor as special service edge, static-network
shall not be advertised if installed as DROP
  timers                   BGP per neighbor timers
```

```
  update-source             Source of routing updates
  peer-dampening            Enable peer-dampening
  peer-group                Assign peer-group to neighbor
NH-Leaf1(config-bgp-ipv4u-vrf)#
```

- Route Specific Configuration

```
NH-Leaf1(config-bgp-ipv4u-vrf)# neighbor 25.x.x.x ?
Possible completions:

  activate                  Allow exchange of route in the current family mode
  additional-paths          Specify bgp additional paths
  allowas-in                Disables the AS_PATH check of the routes learned from
the AS
  capability                Advertise capability to the peer
  default-originate         Originate default route to peer
  enable-peer-as-check      Disable routes advertise between peers in same AS
  filter-list               Establish BGP filters
  maximum-prefix            Maximum number of prefix accept from this peer
  prefix-list               Prefix List for filtering routes
  route-map                 Apply route map to neighbor
  route-reflector-client    Configure a neighbor as Route Reflector client
  send-community            Send community attribute to this neighbor
  unsuppress-map            Route-map to selectively unsuppress suppressed routes
  weight                    Set default weight for routes from this neighbor
NH-Leaf1(config-bgp-ipv4u-vrf)# neighbor 25.x.x.x
```

b. IPv6 Address Family

- Route Specific Configuration

```
NH-Leaf1(config-bgp-ipv6u-vrf)# neighbor 25.x.x.x ?
Possible completions:
  activate                  Allow exchange of route in the current family mode
  additional-paths          Specify bgp additional paths
  allowas-in                Disables the AS_PATH check of the routes learned from
the AS
  capability                Advertise capability to the peer
  default-originate         Originate default route to peer
  enable-peer-as-check      Disable routes advertise between peers in same AS
  filter-list               Establish BGP filters
  maximum-prefix            Maximum number of prefix accept from this peer
  prefix-list               Prefix List for filtering routes
  route-map                 Apply route map to neighbor
  route-reflector-client    Configure a neighbor as Route Reflector client
  send-community            Send community attribute to this neighbor
  unsuppress-map            Route-map to selectively unsuppress suppressed routes
  weight                    Set default weight for routes from this neighbor
NH-Leaf1(config-bgp-ipv6u-vrf)# neighbor 25.x.x.x
```

2. In XCO, configure **neighbor <peer-group> activate** under IPv6 address family of the default VRF when you create a BGP peer group. Multi protocol is enabled by default in Dynamic peers.

```
efa tenant service bgp peer create --name <bgp-service-peer-name> --tenant <tenant-
name> --description <description>
    --ipv4-uc-nbr <device-ip,vrf-name:ipv4-neighbor,remote-as>
    --ipv4-uc-nbr-activate-in-ipv6-uc-af   <device-ip,vrf-name:ipv4-neighbor,true/
false>
    --ipv6-uc-nbr <device-ip,vrf-name:ipv6-neighbor,remote-as>
    --md5-password-prompt-enable=true | false
    --ipv4-uc-nbr-md5-password <device-ip,vrf-name:ipv4-neighbor,ipv4-md5-password>
    --ipv6-uc-nbr-md5-password <device-ip,vrf-name:ipv6-neighbor,ipv6-md5-password>
    --ipv4-uc-nbr-next-hop-self <device-ip,vrf-name:ipv4-neighbor,next-hop-self(true/
false/always)>
    --ipv6-uc-nbr-next-hop-self <device-ip,vrf-name:ipv6-neighbor,next-hop-self(true/
false/always)>
```

```
    --ipv4-uc-nbr-bfd <device-ip,vrf-name:ipv4-neighbor,bfd-enable(true/false),bfd-
interval,bfd-min-rx,bfd-multiplier>
    --ipv6-uc-nbr-bfd <device-ip,vrf-name:ipv6-neighbor,bfd-enable(true/false),bfd-
interval,bfd-min-rx,bfd-multiplier>
    --ipv4-uc-nbr-update-source-ip <device-ip,vrf-name:ipv4-neighbor,update-source-ip>
    --ipv6-uc-nbr-update-source-ip <device-ip,vrf-name:ipv6-neighbor,update-source-ip>
    --ipv4-uc-nbr-remove-private-as <device-ip,vrf-name:ipv4-neighbor,remove-private-
as(true/false)>
    --ipv6-uc-nbr-remove-private-as <device-ip,vrf-name:ipv6-neighbor,remove-private-
as(true/false)>
    --ipv4-uc-nbr-default-originate <device-ip,vrf-name:ipv4-neighbor,default-
originate(true/false)>
    --ipv4-uc-nbr-default-originate-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,default-originate(true/false)>
    --ipv6-uc-nbr-default-originate <device-ip,vrf-name:ipv6-neighbor,default-
originate(true/false)>
    --ipv4-uc-nbr-default-originate-route-map <device-ip,vrf-name:ipv4-neighbor,route-
map-name>
    --ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,route-map-name>
    --ipv6-uc-nbr-default-originate-route-map <device-ip,vrf-name:ipv6-neighbor,route-
map-name>
    --ipv4-uc-nbr-prefix-list <device-ip,vrf-name:ipv4-neighbor,prefix-list-
name,direction(in/out)>
    --ipv4-uc-nbr-prefix-list-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,prefix-
list-name,direction(in/out)>
    --ipv6-uc-nbr-prefix-list <device-ip,vrf-name:ipv6-neighbor,prefix-list-
name,direction(in/out)>
    --ipv4-uc-nbr-route-map <device-ip,vrf-name:ipv4-neighbor,route-map-
name,direction(in/out)>
    --ipv4-uc-nbr-route-map-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,route-map-
name,direction(in/out)>
    --ipv6-uc-nbr-route-map <device-ip,vrf-name:ipv6-neighbor,route-map-
name,direction(in/out)>
    --ipv4-uc-nbr-send-community <device-ip,vrf-name:ipv4-neighbor,all | both | large
| extended | standard | large-and-standard | large-and-extended >
    --ipv4-uc-nbr-send-community-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,all |
both | large | extended | standard | large-and-standard | large-and-extended >
    --ipv6-uc-nbr-send-community <device-ip,vrf-name:ipv6-neighbor,all | both | large
| extended | standard | large-and-standard | large-and-extended >
    --ipv4-uc-nbr-add-path-capability <device-ip,vrf-name:ipv4-neighbor, send |
receive | both>
    --ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,
send | receive | both>
    --ipv6-uc-nbr-add-path-capability <device-ip,vrf-name:ipv6-neighbor, send |
receive | both>
    --ipv4-uc-nbr-add-path-advertise-all <device-ip,vrf-name:ipv4-neighbor,add-path-
advertise-all(true/false)>
    --ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,add-path-advertise-all(true/false)>
    --ipv6-uc-nbr-add-path-advertise-all <device-ip,vrf-name:ipv6-neighbor,add-path-
advertise-all(true/false)>
    --ipv4-uc-nbr-add-path-advertise-best <device-ip,vrf-name:ipv4-neighbor,2-16>
    --ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,2-16>
    --ipv6-uc-nbr-add-path-advertise-best <device-ip,vrf-name:ipv6-neighbor,2-16>
    --ipv4-uc-nbr-add-path-advertise-group-best <device-ip,vrf-name:ipv4-neighbor,add-
path-advertise-group-best(true/false)>
    --ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,add-path-advertise-group-best(true/false)>
    --ipv6-uc-nbr-add-path-advertise-group-best <device-ip,vrf-name:ipv6-neighbor,add-
path-advertise-group-best(true/false)>
    --ipv4-uc-dyn-nbr <device-ip,vrf-name:ipv4-listen-range,peer-group-name,listen-
limit>
```

```
    --ipv6-uc-dyn-nbr <device-ip,vrf-name:ipv6-listen-range,peer-group-name,listen-
limit>

efa tenant service bgp peer update --name <bgp-service-peer-name> --tenant <tenant-
name> --description <description> --operation <peer-add | peer-delete>
    --ipv4-uc-nbr <device-ip,vrf-name:ipv4-neighbor,remote-as>
    --ipv4-uc-nbr-activate-in-ipv6-uc-af   <device-ip,vrf-name:ipv4-neighbor,true/
false>
    --ipv6-uc-nbr <device-ip,vrf-name:ipv6-neighbor,remote-as>
    --md5-password-prompt-enable=true | false
    --ipv4-uc-nbr-md5-password <device-ip,vrf-name:ipv4-neighbor,ipv4-md5-password>
    --ipv6-uc-nbr-md5-password <device-ip,vrf-name:ipv6-neighbor,ipv6-md5-password>
    --ipv4-uc-nbr-next-hop-self <device-ip,vrf-name:ipv4-neighbor,next-hop-self(true/
false/always)>
    --ipv6-uc-nbr-next-hop-self <device-ip,vrf-name:ipv6-neighbor,next-hop-self(true/
false/always)>
    --ipv4-uc-nbr-bfd <device-ip,vrf-name:ipv4-neighbor,bfd-enable(true/false),bfd-
interval,bfd-min-rx,bfd-multiplier>
    --ipv6-uc-nbr-bfd <device-ip,vrf-name:ipv6-neighbor,bfd-enable(true/false),bfd-
interval,bfd-min-rx,bfd-multiplier>
    --ipv4-uc-nbr-update-source-ip <device-ip,vrf-name:ipv4-neighbor,update-source-ip>
    --ipv6-uc-nbr-update-source-ip <device-ip,vrf-name:ipv6-neighbor,update-source-ip>
    --ipv4-uc-nbr-remove-private-as <device-ip,vrf-name:ipv4-neighbor,remove-private-
as(true/false)>
    --ipv6-uc-nbr-remove-private-as <device-ip,vrf-name:ipv6-neighbor,remove-private-
as(true/false)>
    --ipv4-uc-nbr-default-originate <device-ip,vrf-name:ipv4-neighbor,default-
originate(true/false)>
    --ipv4-uc-nbr-default-originate-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,default-originate(true/false)>
    --ipv6-uc-nbr-default-originate <device-ip,vrf-name:ipv6-neighbor,default-
originate(true/false)>
    --ipv4-uc-nbr-default-originate-route-map <device-ip,vrf-name:ipv4-neighbor,route-
map-name>
    --ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,route-map-name>
    --ipv6-uc-nbr-default-originate-route-map <device-ip,vrf-name:ipv6-neighbor,route-
map-name>
    --ipv4-uc-nbr-prefix-list <device-ip,vrf-name:ipv4-neighbor,prefix-list-
name,direction(in/out)>
    --ipv4-uc-nbr-prefix-list-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,prefix-
list-name,direction(in/out)>
    --ipv6-uc-nbr-prefix-list <device-ip,vrf-name:ipv6-neighbor,prefix-list-
name,direction(in/out)>
    --ipv4-uc-nbr-route-map <device-ip,vrf-name:ipv4-neighbor,route-map-
name,direction(in/out)>
    --ipv4-uc-nbr-route-map-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,route-map-
name,direction(in/out)>
    --ipv6-uc-nbr-route-map <device-ip,vrf-name:ipv6-neighbor,route-map-
name,direction(in/out)>
    --ipv4-uc-nbr-send-community <device-ip,vrf-name:ipv4-neighbor,all | both | large
| extended | standard | large-and-standard | large-and-extended >
    --ipv4-uc-nbr-send-community-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,all |
both | large | extended | standard | large-and-standard | large-and-extended >
    --ipv6-uc-nbr-send-community <device-ip,vrf-name:ipv6-neighbor,all | both | large
| extended | standard | large-and-standard | large-and-extended >
    --ipv4-uc-nbr-add-path-capability <device-ip,vrf-name:ipv4-neighbor, send |
receive | both>
    --ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af <device-ip,vrf-name:ipv4-neighbor,
send | receive | both>
    --ipv6-uc-nbr-add-path-capability <device-ip,vrf-name:ipv6-neighbor, send |
receive | both>
    --ipv4-uc-nbr-add-path-advertise-all <device-ip,vrf-name:ipv4-neighbor,add-path-
advertise-all(true/false)>
    --ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
```

```
neighbor,add-path-advertise-all(true/false)>
    --ipv6-uc-nbr-add-path-advertise-all <device-ip,vrf-name:ipv6-neighbor,add-path-
advertise-all(true/false)>
    --ipv4-uc-nbr-add-path-advertise-best <device-ip,vrf-name:ipv4-neighbor,2-16>
    --ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,2-16>
    --ipv6-uc-nbr-add-path-advertise-best <device-ip,vrf-name:ipv6-neighbor,2-16>
    --ipv4-uc-nbr-add-path-advertise-group-best <device-ip,vrf-name:ipv4-neighbor,add-
path-advertise-group-best(true/false)>
    --ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af <device-ip,vrf-name:ipv4-
neighbor,add-path-advertise-group-best(true/false)>
    --ipv6-uc-nbr-add-path-advertise-group-best <device-ip,vrf-name:ipv6-neighbor,add-
path-advertise-group-best(true/false)>
    --ipv4-uc-dyn-nbr <device-ip,vrf-name:ipv4-listen-range,peer-group-name,listen-
limit>
    --ipv6-uc-dyn-nbr <device-ip,vrf-name:ipv6-listen-range,peer-group-name,listen-
limit>
```

Example

```
(efa:root)root@Server41:~# efa tenant service bgp peer create --name customer_1 --tenant
tv3 \
--ipv4-uc-nbr 10.20.49.119,v1:10.10.10.11,95001 \
--ipv4-uc-nbr-md5-password 10.20.49.119,v1:10.10.10.11,password  \
--ipv4-uc-nbr-bfd 10.20.49.119,v1:10.10.10.11,true,50,5000,50 \
--ipv4-uc-nbr-next-hop-self 10.20.49.119,v1:10.10.10.11,always \
--ipv4-uc-nbr-update-source-ip 10.20.49.119,v1:10.10.10.11,10.11.12.13 \
--ipv4-uc-nbr-remove-private-as 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate-route-map 10.20.49.119,v1:10.10.10.11,rt1 \
--ipv4-uc-nbr-route-map 10.20.49.119,v1:10.10.10.11,customer_1_v4_in,in \
--ipv4-uc-nbr-prefix-list 10.20.49.119,v1:10.10.10.11,customer_1_v4_out,out \
--ipv4-uc-nbr-send-community 10.20.49.119,v1:10.10.10.11,all \
--ipv4-uc-nbr-add-path-capability 10.20.49.119,v1:10.10.10.11,both \
--ipv4-uc-nbr-add-path-advertise-all 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-add-path-advertise-best 10.20.49.119,v1:10.10.10.11,16 \
--ipv4-uc-nbr-add-path-advertise-group-best 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-activate-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,rt1 \
--ipv4-uc-nbr-send-community-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,all \
--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,both \
--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,16 \
--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af
10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-route-map-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,customer_1__v6_in,in \
--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,customer_1__v6_out,out
(efa:root)root@Server41:~# efa tenant service bgp peer update --name customer_1 --tenant
tv3 --operation peer-add \
--ipv4-uc-nbr 10.20.49.118,v1:10.10.10.12,95001 \
--ipv4-uc-nbr-md5-password 10.20.49.118,v1:10.10.10.12,password  \
--ipv4-uc-nbr-bfd 10.20.49.118,v1:10.10.10.12,true,50,5000,50 \
--ipv4-uc-nbr-next-hop-self 10.20.49.118,v1:10.10.10.12,always \
--ipv4-uc-nbr-update-source-ip 10.20.49.118,v1:10.10.10.12,10.11.12.12 \
--ipv4-uc-nbr-remove-private-as 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate-route-map 10.20.49.118,v1:10.10.10.12,rt1 \
--ipv4-uc-nbr-route-map 10.20.49.118,v1:10.10.10.12,customer_1_v4_in,in \
--ipv4-uc-nbr-prefix-list 10.20.49.118,v1:10.10.10.12,customer_1_v4_out,out \
--ipv4-uc-nbr-send-community 10.20.49.118,v1:10.10.10.12,all \
--ipv4-uc-nbr-add-path-capability 10.20.49.118,v1:10.10.10.12,both \
--ipv4-uc-nbr-add-path-advertise-all 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-add-path-advertise-best 10.20.49.118,v1:10.10.10.12,16 \
```

```
--ipv4-uc-nbr-add-path-advertise-group-best 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-activate-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,rt1 \
--ipv4-uc-nbr-send-community-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,all \
--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,both \
--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,16 \
--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af
10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-route-map-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,customer_1__v6_in,in \
--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,customer_1__v6_out,out
(efa:root)root@Server41:~# efa tenant service bgp peer show --detail
================================================================================================
=====
Name             : customer_1
Tenant           : tv3
State            : bgp-peer-created
Description      :

Static Peer
-----------
        Device IP                   : 10.20.49.119
        VRF                         : v1
        AFI                         : ipv4
        SAFI                        : unicast
        Remote IP                   : 10.10.10.11
        Remote ASN                  : 95001
        Activate                    : true
        Next Hop Self               : always
        Update Source IP            : 10.11.12.13
        BFD Enabled                 : true
        BFD Interval                : 50
        BFD Rx                      : 5000
        BFD Multiplier              : 50
        MD5 Password                : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS           : true
        Default Originate           : true
        Default Originate Route Map : rt1
        Send Community              : all
        Prefix List In              :
        Prefix List Out             : customer_1__v4_out
        Route Map In                : customer_1__v4_in
        Route Map Out               :
        Add Path Capability         : both
        Add Path Advertise          : All, Group Best, Best 16
        Dev State                   : provisioned
        App State                   : cfg-in-sync
        --Multi protocol capability--
        AFI                         : ipv6
        SAFI                        : unicast
        Activate                    : true
        Default Originate           : true
        Default Originate Route Map : rt1
        Send Community              : all
        Prefix List In              :
        Prefix List Out             : customer_1__v6_out
        Route Map In                : customer_1__v6_in
        Route Map Out               :
        Add Path Capability         : both
        Add Path Advertise          : All, Group Best, Best 16
        Dev State                   : provisioned
        App State                   : cfg-in-sync
```

```
        Device IP                   : 10.20.49.118
        VRF                         : v1
        AFI                         : ipv4
        SAFI                        : unicast
        Remote IP                   : 10.10.10.12
        Remote ASN                  : 95001
        Activate                    : true
        Next Hop Self               : always
        Update Source IP            : 10.11.12.12
        BFD Enabled                 : true
        BFD Interval                : 50
        BFD Rx                      : 5000
        BFD Multiplier              : 50
        MD5 Password                : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS           : true
        Default Originate           : true
        Default Originate Route Map : rt1
        Send Community              : all
        Prefix List In              :
        Prefix List Out             : customer_1__v4_out
        Route Map In                : customer_1__v4_in
        Route Map Out               :
        Add Path Capability         : both
        Add Path Advertise          : All, Group Best, Best 16
        Dev State                   : provisioned
        App State                   : cfg-in-sync
        --Multi protocol capability--
        AFI                         : ipv6
        SAFI                        : unicast
        Activate                    : true
        Default Originate           : true
        Default Originate Route Map : rt1
        Send Community              : all
        Prefix List In              :
        Prefix List Out             : customer_1__v6_out
        Route Map In                : customer_1__v6_in
        Route Map Out               :
        Add Path Capability         : both
        Add Path Advertise          : All, Group Best, Best 16
        Dev State                   : provisioned
        App State                   : cfg-in-sync


Dynamic Peer
------------
        0 Records
```

```
================================================================================
====
```

```
Rack1-Device1# show runn router bgp
router bgp
 local-as 65000
 capability as4-enable
 fast-external-fallover
 neighbor podA-spine-group peer-
group
 neighbor podA-spine-group remote-
as 64512
 neighbor podA-spine-group
description To Spine
 neighbor podA-spine-
group password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor podA-spine-group bfd
 neighbor 10.10.10.36 peer-group
podA-spine-group
 neighbor 10.20.20.11 remote-as
65000
 neighbor 10.20.20.11 next-hop-self
 neighbor 10.20.20.11 password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor 10.20.20.11 bfd
 address-family ipv4 unicast
  network 172.31.254.170/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf v1
  redistribute connected
  additional-paths select all
  neighbor 10.10.10.11 remote-as
95001
  neighbor 10.10.10.11 default-
originate route-map rt1
  neighbor 10.10.10.11 prefix-list
customer_1__v4_out out
  neighbor 10.10.10.11 route-map
in customer_1__v4_in
  neighbor 10.10.10.11 additional-
paths send receive
  neighbor 10.10.10.11 additional-
paths advertise best 16 group-best
all
  neighbor 10.10.10.11 next-hop-
self always
  neighbor 10.10.10.11 password
$9$MCgKGaNt6OASX68/7TC6Lw==
  neighbor 10.10.10.11 remove-
private-as
  neighbor 10.10.10.11 bfd
  neighbor 10.10.10.11 bfd
interval 50 min-rx 5000 multiplier
50
  neighbor 10.10.10.11 update-
source 10.11.12.13
  neighbor 10.10.10.11 send-
community all
  neighbor 10.40.40.250 remote-as
65000
```

```
Rack1-Device2# show runn router bgp
router bgp
 local-as 65000
 capability as4-enable
 fast-external-fallover
 neighbor podA-spine-group peer-
group
 neighbor podA-spine-group remote-
as 64512
 neighbor podA-spine-group
description To Spine
 neighbor podA-spine-
group password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor podA-spine-group bfd
 neighbor 10.10.10.33 peer-group
podA-spine-group
 neighbor 10.20.20.10 remote-as
65000
 neighbor 10.20.20.10 next-hop-self
 neighbor 10.20.20.10 password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor 10.20.20.10 bfd
 address-family ipv4 unicast
  network 172.31.254.170/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf v1
  redistribute connected
  additional-paths select all
  neighbor 10.10.10.12 remote-as
95001
  neighbor 10.10.10.12 default-
originate route-map rt1
  neighbor 10.10.10.12 prefix-list
customer_1__v4_out out
  neighbor 10.10.10.12 route-map
in customer_1__v4_in
  neighbor 10.10.10.12 additional-
paths send receive
  neighbor 10.10.10.12 additional-
paths advertise best 16 group-best
all
  neighbor 10.10.10.12 next-hop-
self always
  neighbor 10.10.10.12 password
$9$MCgKGaNt6OASX68/7TC6Lw==
  neighbor 10.10.10.12 remove-
private-as
  neighbor 10.10.10.12 bfd
  neighbor 10.10.10.12 bfd
interval 50 min-rx 5000 multiplier
50
  neighbor 10.10.10.12 update-
source 10.11.12.12
  neighbor 10.10.10.12 send-
community all
  neighbor 10.40.40.251 remote-as
65000
```

```
  neighbor 10.40.40.250 next-hop-
self
  neighbor 10.40.40.250 password
$9$MCgKGaNt6OASX68/7TC6Lw==
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf v1
  redistribute connected
  additional-paths select all
  neighbor fd40:4040:4040:1::fc
remote-as 65000
  neighbor fd40:4040:4040:1::fc
next-hop-self
  neighbor
fd40:4040:4040:1::fc password
$9$MCgKGaNt6OASX68/7TC6Lw==
  neighbor fd40:4040:4040:1::fc
activate
  neighbor 10.10.10.11 activate
  neighbor 10.10.10.11 send-
community all
  neighbor 10.10.10.11 default-
originate route-map rt1
  neighbor 10.10.10.11 prefix-list
customer_1__v6_out out
  neighbor 10.10.10.11 route-map
in customer_1__v6_in
  neighbor 10.10.10.11 additional-
paths send receive
  neighbor 10.10.10.11 additional-
paths advertise best 16 group-best
all
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
  neighbor podA-spine-group
encapsulation vxlan
  neighbor podA-spine-group next-
hop-unchanged
  neighbor podA-spine-group enable-
peer-as-check
  neighbor podA-spine-group
activate
 !
!
Rack1-Device1#
```

```
  neighbor 10.40.40.251 next-hop-
self
  neighbor 10.40.40.251 password
$9$MCgKGaNt6OASX68/7TC6Lw==
  maximum-paths 8
 !
 address-family ipv6 unicast
 !
 address-family ipv6 unicast vrf v1
  redistribute connected
  additional-paths select all
  neighbor fd40:4040:4040:1::fd
remote-as 65000
  neighbor fd40:4040:4040:1::fd
next-hop-self
  neighbor
fd40:4040:4040:1::fd password
$9$MCgKGaNt6OASX68/7TC6Lw==
  neighbor fd40:4040:4040:1::fd
activate
  neighbor 10.10.10.12 activate
  neighbor 10.10.10.12 send-
community all
  neighbor 10.10.10.12 default-
originate route-map rt1
  neighbor 10.10.10.12 prefix-list
customer_1__v6_out out
  neighbor 10.10.10.12 route-map
in customer_1__v6_in
  neighbor 10.10.10.12 additional-
paths send receive
  neighbor 10.10.10.12 additional-
paths advertise best 16 group-best
all
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
  neighbor podA-spine-group
encapsulation vxlan
  neighbor podA-spine-group next-
hop-unchanged
  neighbor podA-spine-group enable-
peer-as-check
  neighbor podA-spine-group
activate
 !
!
Rack1-Device2#
```

Configure Multi Protocol BGP on Tenant Dynamic BGP Peer

You can configure multi protocol BGP.

About This Task

Follow this procedure to configure multi protocol BGP on tenant dynamic BGP peer.

Procedure

1. On SLX devices, run the following command to activate an associated BGP peer group under the IPv6 address family of a default VRF:

```
router bgp
  local-as 100
  neighbor pg1 peer-group
  neighbor pg1 remote-as 100
  listen-range 10.1.0.0/16 peer-group pg1 limit 20
  address-family ipv6 unicast
    neighbor pg1 activate
  !
!
```

2. In XCO, configure **neighbor <peer-group> activate** under IPv6 address family of a default VRF when you create a BGP peer group. Multi protocol is enabled by default in Dynamic peers.

Enable or Disable MP BGP Capability for IPv6 Prefix Exchange over IPv4 Peer

You can enable or disable multi protocol BGP on tenant BGP peer group by adding IPv6 prefix over the IPv4 peer for all the BGP peer groups.

About This Task

Follow this procedure to enable or disable IPv6 prefix over IPv4 peer.

Note

- The IPv6 Prefix over IPv4 Peer feature for all the BGP peer groups is disabled by default on SLX.
- Ensure the following configuration changes to enable the IPv4 BGP peers associated with the BGP peer groups to carry the IPv6 prefixes:
  - Enable the IPv6 Prefix Over IPv4 Peer feature.
  - Activate the BGP peer group under the IPv6 address-family of the default-vrf.
- Upgrading XCO to the version 3.3.1 and later does not impact existing peer groups and peers. The configuration stays intact.
- The IPv4 BGP sessions mapped to the peer group with IPv6 address family enabled will be cleared when the IPv6-Prefix over IPv4-Peer is enabled or disabled.

Procedure

Run the following command to enable or disable IPv6 prefix over IPv4 peer for all the BGP peer groups in the device inventory:

```
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 OS ONE device 30s-24h
Example. 30s or 99m or 1h20m or 20m, default 6m for Fabric device, 30s for OS ONE 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. This configuration is applicable for Extreme 8520, 8720 and
8820 hardware
      --maximum-ipv6-prefix-length-64-urpf string      Enter Yes to configure the maximum
route prefix length of 64 along with unicast reverse path forwarding. This configuration
is applicable for Extreme 8520, 8720 and 8820 hardware
      --peer-group-ipv6-prefix-over-ipv4-peer string   Enter Yes to enable the peer group
ipv6prefix over ipv4peer. Enter No to disable. This configuration is supported from the
firmware version 20.5.2a. IPv4 BGP sessions mapped to peer group with IPv6 address family
enabled will be cleared.
```

The following is an example of enabling or disabling IPv6 prefix over IPv4 peer for all the BGP peer groups in the system:

```
efa inventory device setting update --ip 10.20.x.x --peer-group-ipv6-prefix-over-ipv4-
peer yes
+--------------+----------------------------------------+---------+-------+-------+
|  IP ADDRESS  |                NAME                    | STATUS  | VALUE | ERROR |
+--------------+----------------------------------------+---------+-------+-------+
| 10.20.x.x    | peer-group-ipv6-prefix-over-ipv4-peer | Success | Yes   |       |
+--------------+----------------------------------------+---------+-------+-------+


efa inventory device setting show --ip 10.20.x.x
+--------------------------------+-------+-------------+
|              NAME              | VALUE |  APP STATE  |
+--------------------------------+-------+-------------+
| 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    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Static Prefix Independent      | No    |             |
| Convergence                    |       |             |
+--------------------------------+-------+-------------+
| Maximum Load Sharing Paths     |       |             |
+--------------------------------+-------+-------------+
| Maximum Ipv6 Prefix Length 64  |       |             |
+--------------------------------+-------+-------------+
| Urpf                           |       |             |
+--------------------------------+-------+-------------+
| Ip Option Disable              | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Ip Option Disable Cpu          | No    |             |
+--------------------------------+-------+-------------+
| Ipv6 Option Disable            | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Peer Group Ipv6 Prefix Over    | Yes   | cfg-in-sync |
| Ipv4 Peer                      |       |             |
+--------------------------------+-------+-------------+
Rack1-Device1# sh run router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 peer-group ipv6prefix-over-ipv4peer
 fast-external-fallover
 neighbor 10.20.y.y remote-as 4200000000
 neighbor 10.20.y.y next-hop-self
 address-family ipv4 unicast
  network 172.xx.254.119/32
  network 172.xx.254.215/32
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
```

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 520
2. Configure IP Prefix List and Route Map on Tenant BGP Peer Group on page 521
3. Configure Send-Community on Tenant BGP Peer Group on page 525

4. Add Path on Tenant BGP Peer Group on page 528
5. Configure remove-private-as on BGP Peer Group on page 532
6. Activate Peer Group on Tenant BGP on page 534
7. Delete Pending BGP Peer Group Configuration on page 537
8. Configure IPv6 Address as Update Source on page 538

Create a BGP Peer Group

You can configure a BGP peer group.

About This Task

Follow this procedure to configure a BGP peer group.

Note

For detailed information on flag requirements and component relationships during tenant and its components updates, see Tenant Components Dependencies and Creation Workflow and Flag Requirements for Tenant Components Update Operations.

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>
    --pg-ipv6-uc-nbr-activate <device-ip,pg-name:true/false>
```

   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
    --pg-ipv6-uc-nbr-activate 10.20.246.29,v1:true
```

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>
    --pg-ipv6-uc-nbr-activate 10.20.246.29,v1:true
```

   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 --pg-ipv6-uc-
nbr-activate 10.20.246.29,v1:true
```

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, 4.1.0*

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, 4.1.0*.

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 does not activate the BGP peer group when you create a BGP peer group.

Note

- During upgrade, all the BGP peer groups without an IPv6 listen-range association are deactivated from the IPv6 address-family.
- An upgrade from XCO 3.2.1 to XCO 3.3.0 and later will refresh all the existing peer groups which are not associated with any IPv6 listen-range.
- When you upgrade from XCO 3.2.1 to XCO 3.3.0 and later, all the existing peer groups, which are not associated with any IPv6 listen-range, will be refreshed. Within 30 min from the upgrade, an auto DRC gets initiated,and deactivates all such Peer-groups under IPv6 address-family, and moves them back to the in-sync state.
- Be aware that it takes 15 minutes to get it reflected in device and XCO, and till then XCO reflects it in the drift state. As a best practice, do not perform a manual DRC.

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 --tenant "t1" --name "vs"
         --pg-name 10.20.246.30:v1 --pg-asn 10.20.246.30,v1:5200
         --pg-ipv6-uc-nbr-activate 10.20.246.30,v1:true
         --pg-ipv4-uc-nbr-route-map 10.20.246.30,v1:customer_1_in,in
         --pg-ipv4-uc-nbr-prefix-list 10.20.246.30,v1:customer_1_in,in
         --pg-name 10.20.246.30:v2 --pg-asn 10.20.246.30,v2:5201
         --pg-name 10.20.246.30:v3 --pg-asn 10.20.246.30,v3:5203
         --pg-name 10.20.246.30:v4 --pg-asn 10.20.246.30,v4:5204

         --pg-name 10.20.246.29:v1 --pg-asn 10.20.246.29,v1:5200
         --pg-ipv6-uc-nbr-activate 10.20.246.29,v1:true
         --pg-ipv4-uc-nbr-route-map 10.20.246.29,v1:customer_1_in,in
         --pg-ipv4-uc-nbr-prefix-list 10.20.246.29,v1:customer_1_in,in
         --pg-name 10.20.246.29:v2 --pg-asn 10.20.246.29,v2:5201
         --pg-name 10.20.246.29:v3 --pg-asn 10.20.246.29,v3:5203
         --pg-name 10.20.246.29:v4 --pg-asn 10.20.246.29,v4:5204
efa tenant service bgp peer create --tenant "t1" --name "vs"
         --ipv4-uc-dyn-nbr 10.20.246.29,vs:15.16.16.0/28,v1,20
         --ipv4-uc-dyn-nbr 10.20.246.29,vs:15.16.17.0/28,v3,20
         --ipv4-uc-dyn-nbr 10.20.246.29,vs:15.16.18.0/28,v4,20
         --ipv6-uc-dyn-nbr 10.20.246.29,vs:14::/127,v2,10

         --ipv6-uc-dyn-nbr 10.20.246.29,vs:15::/127,v3,10
         --ipv4-uc-dyn-nbr 10.20.246.30,vs:15.16.16.0/28,v1,20
         --ipv4-uc-dyn-nbr 10.20.246.30,vs:15.16.17.0/28,v3,20
         --ipv4-uc-dyn-nbr 10.20.246.30,vs:15.16.18.0/28,v4,20
```

```
        --ipv6-uc-dyn-nbr 10.20.246.30,vs:14::/127,v2,10
        --ipv6-uc-dyn-nbr 10.20.246.30,vs:15::/127,v3,10
```

```
Rack1-Device1# show runn router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor v1 peer-group
 neighbor v1 remote-as 5200
 neighbor v2 peer-group
 neighbor v2 remote-as 5201
 neighbor v3 peer-group
 neighbor v3 remote-as 5203
 neighbor v4 peer-group
 neighbor v4 remote-as 5204
  neighbor  10.20.20.7  remote-as
4200000000
 neighbor 10.20.20.7 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.35/32
  network 172.31.254.185/32
      neighbor   v1   prefix-list
customer_1_in in
     neighbor   v1   route-map   in
customer_1_in
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf vs
  redistribute connected
   listen-range 15.16.16.0/28 peer-
group v1 limit 20
   listen-range 15.16.17.0/28 peer-
group v3 limit 20
   listen-range 15.16.18.0/28 peer-
group v4 limit 20
  maximum-paths 8
 !
 address-family ipv6 unicast
  neighbor v1 activate
  neighbor v2 activate
  neighbor v3 activate
 !
 address-family ipv6 unicast vrf vs
  redistribute connected
  listen-range 14::/127 peer-group
v2 limit 10
  listen-range 15::/127 peer-group
v3 limit 10
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
Rack1-Device1#
```

```
Rack1-Device1# show runn router bgp
router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor v1 peer-group
 neighbor v1 remote-as 5200
 neighbor v2 peer-group
 neighbor v2 remote-as 5201
 neighbor v3 peer-group
 neighbor v3 remote-as 5203
 neighbor v4 peer-group
 neighbor v4 remote-as 5204
  neighbor  10.20.20.6  remote-as
4200000000
 neighbor 10.20.20.6 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.35/32
  network 172.31.254.66/32
      neighbor   v1   prefix-list
customer_1_in in
     neighbor   v1   route-map   in
customer_1_in
  maximum-paths 8
  graceful-restart
 !
 address-family ipv4 unicast vrf vs
  redistribute connected
   listen-range 15.16.16.0/28 peer-
group v1 limit 20
   listen-range 15.16.17.0/28 peer-
group v3 limit 20
   listen-range 15.16.18.0/28 peer-
group v4 limit 20
  maximum-paths 8
 !
 address-family ipv6 unicast
  neighbor v1 activate
  neighbor v2 activate
  neighbor v3 activate
 !
 address-family ipv6 unicast vrf vs
  redistribute connected
  listen-range 14::/127 peer-group
v2 limit 10
  listen-range 15::/127 peer-group
v3 limit 10
  maximum-paths 8
 !
 address-family l2vpn evpn
  graceful-restart
 !
!
Rack1-Device1#
```

Delete Pending BGP Peer Group Configuration

You can delete pending configuration on a BGP peer group.

About This Task

Follow this procedure to remove the pending configuration on a BGP peer group.

Procedure

Run the following command:

```
efa tenant service bgp peer-group configure
```

The **`efa tenant service bgp peer-group configure`** command pushes or removes a pending configuration on a BGP peer group instance when it is in one of the following states:

`bgp-pg-delete-pending | bgp-pg-peer-group-delete-pending | bgp-pg-peer-group-activate-pending`

Example

```
efa tenant service bgp peer-group show
================================================================================================
=====================================================================
Name        : tv3_pg1
Tenant      : tv3
State       : bgp-pg-peer-group-delete-pending

+-------------+------------+--------+----------------------+----------+------------
+---------+-------------------------+-------------+-------------+
|  Device IP  | Peer Group | Remote |        Activate        | Next Hop |   Update   |
BFD   |           BFD            |  Dev State  |  App State  |
|             |            |  ASN   |  [AFI,SAFI->Activate]  |   Self   | Source IP  |
Enabled | [Interval,Rx,Multiplier] |             |             |
+-------------+------------+--------+----------------------+----------+------------
+---------+-------------------------+-------------+-------------+
| 10.20.61.91 |    pg1     | 95002  | ipv4, unicast -> true  |   true   | 10.10.10.3 |
true   |       660, 506, 20       | provisioned | cfg-in-sync |
|             |            |        | ipv6, unicast -> false |          |
|         |                          |             |             |
+-------------+------------+--------+----------------------+----------+------------
+---------+-------------------------+-------------+-------------+
| 10.20.61.90 |    pg1     | 95002  | ipv4, unicast -> true  |   true   | 10.10.10.3 |
true   |       660, 506, 20       | provisioned | cfg-in-sync |
|             |            |        | ipv6, unicast -> false |          |
|         |                          |             |             |
+-------------+------------+--------+----------------------+----------+------------
+---------+-------------------------+-------------+-------------+
BGP PeerGroup Details

================================================================================================
=====================================================================

--- Time Elapsed: 433.275989ms ---


(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer-group configure --name tv3_pg1
--tenant tv3

BgpService configured successfully.

--- Time Elapsed: 8.920535933s ---
```

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer-group show
================================================================================================
================================================================================
Name        : tv3_pg1
Tenant      : tv3
State       : bgp-pg-created

+-------------+------------+--------+-----------------------+----------+------------
+---------+--------------------------+-------------+-------------+
|  Device IP  | Peer Group | Remote |        Activate       | Next Hop |   Update   |
BFD   |           BFD            |  Dev State  |  App State  |
|             |            |  ASN   |  [AFI,SAFI->Activate] |   Self   | Source IP  |
Enabled | [Interval,Rx,Multiplier] |             |             |
+-------------+------------+--------+-----------------------+----------+------------
+---------+--------------------------+-------------+-------------+
| 10.20.61.91 |    pg1     | 95002  | ipv4, unicast -> true |   true   | 10.10.10.3 |
true   |       660, 506, 20       | provisioned | cfg-in-sync |
|             |            |        | ipv6, unicast -> false |          |
|         |                          |             |             |
+-------------+------------+--------+-----------------------+----------+------------
+---------+--------------------------+-------------+-------------+
BGP PeerGroup Details

================================================================================================
================================================================================

--- Time Elapsed: 433.275989ms ---
```

Configure IPv6 Address as Update Source

You can configure IPv6 address as an update-source.

About This Task

Follow this procedure to configure IPv6 address as an update-source in the XCO fabric.

Starting from XCO 3.4.0 and SLX 20.5.30 and later, tenant BGP peer-groups support both IPv4 and IPv6 addresses as an update-source. The feature is supported on the following platforms: SLX 9540, SLX 9640, SLX 9740, Extreme 8520, Extreme 8820, and Extreme 8720.

Note

- XCO version 3.4.0 does not support Dual IPv6 addresses, which is IPv6 plus IPv4 formats.
- XCO versions below 3.4.0 and SLX versions below 20.5.3 support update-source attribute for tenant BGP peer-groups, but only IPv4 address can be configured as update-source.

Procedure

1. Run the following command to configure IPv6 address as an update-source when you create a BGP peer-group:

```
efa tenant service bgp peer-group create --tenant <tenant-name> --name <peer-group-
name>
               --pg-name device-ip:peer-group-name --pg-update-source-ip device-
ip,peer-group-name:update-source-ip
```

2. Run the following command to configure IPv6 address as an update-source when you update a BGP peer-group:

```
efa tenant service bgp peer-group update --tenant <tenant-name> --name <peer-group-
name> --operation peer-group-add
              --pg-name device-ip:peer-group-name --pg-update-source-ip device-
ip,peer-group-name:update-source-ip
```

Example

```
efa tenant service bgp peer-group create --tenant "tenant1" --name "bgpPeerGroup2"
            --pg-name 10.20.246.15:bgppg2 --pg-asn 10.20.246.15,bgppg2:100 --pg-update-
source-ip 10.20.246.15,bgppg2:10::10
efa tenant service bgp peer-group update --tenant "tenant1" --name "bgpPeerGroup2" --
operation peer-group-add
            --pg-name 10.20.246.15:bgppg1 --pg-asn 10.20.246.15,bgppg1:200 --pg-update-
source-ip 10.20.246.15,bgppg1:10::20

efa tenant service bgp peer-group show --tenant tenant1 -detail
================================================================================
Name             : bgpPeerGroup2
Tenant           : tenant1
State            : bgp-pg-created
Description      :

Peer Group
----------
        Device IP            : 10.20.246.15
        Peer Group           : bgppg2
        Remote ASN           : 100
        Next Hop Self        : false
        Update Source IP     : 10::10
        BFD Enabled          : false
        BFD Interval         :
        BFD Rx               :
        BFD Multiplier       :
        MD5 Password         :
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  :
        Add Path Advertise   :
        Dev State            : provisioned
        App State            : cfg-in-sync

        Device IP            : 10.20.246.15
        Peer Group           : bgppg1
        Remote ASN           : 200
        Next Hop Self        : false
        Update Source IP     : 10::20
        BFD Enabled          : false
        BFD Interval         :
        BFD Rx               :
        BFD Multiplier       :
        MD5 Password         :
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
```

```
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  :
        Add Path Advertise   :
        Dev State            : provisioned
        App State            : cfg-in-sync
================================================================================
```

```
Rack1-Device1# sh run router bgp
router bgp
local-as 4200000000
capability as4-enable
fast-external-fallover
neighbor bgppg2 peer-group
neighbor bgppg2 remote-as 100
neighbor bgppg2 update-source 10::10
neighbor bgppg1 peer-group
neighbor bgppg1 remote-as 200
neighbor bgppg1 update-source 10::20
neighbor 10.20.20.3 remote-as 4200000000
neighbor 10.20.20.3 next-hop-self
address-family ipv4 unicast
  network 172.31.254.24/32
  network 172.31.254.43/32
  maximum-paths 8
  graceful-restart
!
address-family ipv6 unicast
!
address-family l2vpn evpn
  graceful-restart
!
!Rack1-Device1#
```

BGP Peer Groups with Route Map Policies on OS ONE Devices

Learn about the support for associating route-maps with BGP peer-groups on OS ONE devices through XCO.

You can configure import and export routing policies under BGP peer-group address families for both IPv4 and IPv6 unicast environments. The implementation includes CLI-based configuration workflows and device synchronization support.

XCO supports route-map association for BGP peer-groups on OS ONE devices with the following capabilities:

- Route-map association under BGP peer-groups
- IPv4 Unicast policy support
- IPv6 Unicast policy support
- Import and export policy configuration
- CRUD operations through XCO CLI
- Device synchronization support

Supported Platforms: SLX-OS and OS ONE

Supported Address Families

The solution supports the following address families:

Address Family	Direction Support
IPv4 Unicast	Inbound and Outbound
IPv6 Unicast	Inbound and Outbound

Supported Operations

The following operations are supported for BGP peer-group route-map configuration:

- Create
- Read
- Update
- Delete
- Device synchronization

BGP Peer-Group Route-Map Association

XCO enables route-map association under the BGP peer-group address-family configuration.

Supported Policy Directions

- Import policy (Inbound)
- Export policy (Outbound)

Supported Configuration Scope

- IPv4 Unicast
- IPv6 Unicast

Platform Considerations: OS ONE vs SLX-OS

OS ONE and SLX-OS devices have important behavioral differences in BGP peer group configuration.

Feature	OS ONE	SLX-OS	Implication
Route Map Support	At peer group level only	At peer group level	OS ONE does not support route maps directly on peers
Prefix List Support	Not supported	At peer group level	Use route maps instead of prefix lists on OS ONE
Multiple Policies	Multiple per direction	One per direction	OS ONE allows multiple import/ export policies

Feature	OS ONE	SLX-OS	Implication
Policy Association	Via route maps only	Via route maps or prefix lists	OS ONE simplifies with route maps only
VRF Support	Full VRF support	Default VRF only	Explicit VRF specification in OS ONE

Important Limitations for OS ONE

Be aware of the following limitations when configuring OS ONE devices:

- Prefix lists cannot be associated directly with peer groups – use route maps instead
- Route maps cannot be configured at the individual peer level – only at peer group level

Unsupported Features

The following features are not supported on OS ONE devices.

- Peer-Level Route-Map Association

  The following configuration is not supported:

  - Prefix-list association under BGP peer
  - Route-map association under BGP peer

  Note
  OSONE devices support route-map and prefix-list policies only at the peer-group level.

- Independent Prefix-List Association Under Peer-Group

  Independent prefix-list configuration under a peer-group is not supported.

  Prefix-lists must be associated indirectly through route-maps.

Configure BGP Peer Groups on OS ONE Devices

You can configure BGP peer groups on OS ONE devices.

Before You Begin

Before configuring BGP peer groups, ensure:

- Tenant is created in XCO
- OS ONE device is registered and assigned to fabric
- BGP is enabled on the device
- Required VRF exists on the device
- Route maps have been created (if using route map policies)

About This Task
Follow this procedure to configure BGP peer groups on OS ONE devices.

Procedure

1. Before creating BGP peer groups, prepare the necessary policy objects.
   a. Create route maps.

      Route maps are the policy engine for peer groups on OS ONE. Create route maps before peer group configuration.

      ```
      efa policy route-map create --name <route-map-name> --rule
      seq[<sequence>],action[<permit/deny>]
      ```

      For example, create a route map that permits routes.

      ```
      efa policy route-map create --name rm_import_policy --rule seq[1],action[permit]
      ```

   b. Associate Route Maps to Device.

      After creating route maps, add them to the OS ONE device:

      ```
      efa policy route-map update --name <route-map-name> --operation add-device --ip
      <device-ip>
      ```

2. Create a BGP peer group with route map policies applied. Peer groups must include a name, tenant, peer group name, ASN, and optional route map associations.

   ```
   efa tenant service bgp peer-group create --name <service-name> --tenant <tenant-
   name> --pg-name <device-ip>:<peer-group-name>:<vrf-name> --pg-asn <device-ip>,<peer-
   group-name>,<vrf-name>:<remote-asn> --pg-ipv4-uc-nbr-route-map <device-ip>,<peer-group-
   name>,<vrf-name>:<route-map>,<direction>
   ```

   - The following example creates an IPv4 peer group with import and export route maps:

     ```
     efa tenant service bgp peer-group create --name bgp_service_ipv4 --tenant
     customer1 --pg-name 10.64.184.86:pg1:vrf1 --pg-asn 10.64.184.86,pg1,vrf1:65001 --pg-
     ipv4-uc-nbr-route-map 10.64.184.86,pg1,vrf1:rm_import,in --pg-ipv4-uc-nbr-route-map
     10.64.184.86,pg1,vrf1:rm_export,out
     ```

   - The following example creates an IPv6 peer group with dual-direction route maps:

     ```
     efa tenant service bgp peer-group create --name bgp_service_ipv6 --tenant
     customer1 --pg-name 10.64.184.86:pg2:vrf1 --pg-asn 10.64.184.86,pg2,vrf1:65002 --
     pg-ipv6-uc-nbr-route-map 10.64.184.86,pg2,vrf1:rm6_import,in --pg-ipv6-uc-nbr-route-
     map 10.64.184.86,pg2,vrf1:rm6_export,out
     ```

3. View peer group configuration.

   Display all peer groups in a tenant:

   ```
   efa tenant service bgp peer-group show --tenant <tenant-name>
   ```

   View detailed peer group information:

   ```
   efa tenant service bgp peer-group show --tenant <tenant-name> --detail
   ```

4. Update peer group configuration.

   Modify an existing peer group, such as adding a new peer group to the same service:

   ```
   efa tenant service bgp peer-group update --operation peer-group-add --name <service-
   name> --tenant <tenant-name> --pg-name <device-ip>:<pg-name>:<vrf> --pg-asn <device-
   ip>,<pg-name>,<vrf>:<remote-asn>
   ```

5. Delete peer group.

   Remove a peer group from the service:

   ```
   efa tenant service bgp peer-group delete --name <service-name> --tenant <tenant-name>
   ```

Configure BGP Peers on OS ONE Devices

You can configure BGP peers on OS ONE devices.

About This Task
Follow this procedure to configure BGP peers on OS ONE devices.

Procedure

1. Create a BGP peer.

   Associate a neighbor with a peer group:

   ```
   efa tenant service bgp peer create --name <peer-name> --tenant <tenant-name> --ipv4-uc-nbr <device-ip>,<vrf>:<neighbor-ip>,<peer-group-name>
   ```

   The following example creates IPv4 peer referencing peer group:

   ```
   efa tenant service bgp peer create --name uplink_neighbor_1 --tenant customer1 --ipv4-uc-nbr 10.64.184.86,vrf1:192.168.1.1,pg1
   ```

   This peer inherits all policies (route maps) from peer group pg1.
   The following example creates IPv6 peer referencing peer group:

   ```
   efa tenant service bgp peer create --name uplink_neighbor_v6 --tenant customer1 --ipv6-uc-nbr 10.64.184.86,vrf1:2001:db8::1,pg2
   ```

2. View BGP peers.

   ```
   efa tenant service bgp peer show --tenant <tenant-name>
   ```

   View detailed peer information including policy inheritance:

   ```
   efa tenant service bgp peer show --tenant <tenant-name> --detail
   ```

Configure Route Maps on OS ONE Devices

You can configure route maps on OS ONE devices.

About This Task
Follow this procedure to configure route maps on OS ONE devices.

Procedure

1. Create route maps for BGP.

   Create a basic route map that permits all routes:

   ```
   efa policy route-map create --name allow_all --rule seq[1],action[permit]
   ```

   Create a route map that denies specific routes:

   ```
   efa policy route-map create --name deny_private --rule seq[1],action[deny]
   ```

2. Apply route maps to device.

   After creating route maps, add them to the target device:

   ```
   efa policy route-map update --name <route-map-name> --operation add-device --ip <device-ip>
   ```

3. Using route maps with peer groups.

   Specify route maps in two directions:

   - Import policy (in) – Controls which routes are accepted from neighbors
   - Export policy (out) – Controls which routes are advertised to neighbors

Common Configuration Patterns: Configure Customer BGP Session with Import/Export Control

You can configure a peer group for customer BGP session with inbound and outbound filtering.

About This Task

Follow this procedure to configure a peer group for customer BGP session with inbound and outbound filtering.

Procedure

1. Create import and export route maps.

```
efa policy route-map create --name customer_import --rule seq[1],action[permit]
efa policy route-map create --name customer_export --rule seq[1],action[permit]
```

2. Add route maps to device.

```
efa policy route-map update --name customer_import --operation add-device --ip
10.64.184.86
efa policy route-map update --name customer_export --operation add-device --ip
10.64.184.86
```

3. Create peer group with route maps.

```
efa tenant service bgp peer-group create --name cust_bgp_service --tenant t1 --pg-
name 10.64.184.86:cust_pg:vrf1 --pg-asn 10.64.184.86,cust_pg,vrf1:65001 --pg-ipv4-uc-
nbr-route-map 10.64.184.86,cust_pg,vrf1:customer_import,in --pg-ipv4-uc-nbr-route-map
10.64.184.86,cust_pg,vrf1:customer_export,out
```

4. Create peers referencing the peer group.

```
efa tenant service bgp peer create --name cust_uplink --tenant t1 --ipv4-uc-nbr
10.64.184.86,vrf1:203.0.113.1,cust_pg
```

Tenant Components Dependencies and Creation Workflow

Use this workflow to learn about the dependencies and creation order of tenant components.

Workflow Sequence	Description
Device Registration & Port Discovery	• Register device in Inventory • Inventory fetches device ports and LLDP data
Port Breakout Management	• Breakout ports allowed post-registration and pre-tenant usage • Restrictions: No breakout changes (set or unset breakout) after tenant port consumption
Port Publication to Tenant	• Publish port list to Tenant service post-Fabric config • Create Tenant using published ports
Tenant-Level Port Usage	• Create Port-Channel (PO) using same ports post-Tenant creation • Use PO in L2 and L3 EPG creation
L3 EPG & VRF Relationship	• Create VRF under the tenant, then use the VRF to create L3 EPG

Workflow Sequence	Description
Peer-Group Creation	• Extreme OS ONE: Create Peer-Group under user VRF post-L3 EPG configuration using that VRF • SLX-OS: Create Peer-Group under default-vrf post-tenant creation without dependency on L3 EPG
Static & Dynamic Peers	• Create Static and Dynamic Peers only on a VRF (requires L3 EPG with targeted VRF) • Dynamic Peers: Create on Peer-Groups only (requires Peer-Group first) • Static Peers: ◦ SLX-OS: Create without Peer-Group, immediately after L3 EPG creation ◦ Extreme OS ONE: Create Peer-Group first, then Static Peers

Flag Requirements for Tenant Components Update Operations

Use this topic to learn about the summary of flags needed for updating Tenant, PO, VRF, EPG, Peer-Group, and Peer.

Common Requirements for All Update Operations

To identify an update operation, the following **three flags are always required**:

- `--name`: Identifies the entity being updated
- `--tenant`: Specifies the tenant under which the entity resides
- `--operation`: Specifies which update action is being performed

Tenant Update Operations

Operation	Flags	Notes
desc-update	`--description`	Mandatory
vni-update	`--l2-vni-range, --l3-vni-range`	Either or both can be provided
port-add / port-delete	`--port`	Mandatory
vlan-add	`--vlan-range`	For adding VLAN or VLAN-range
vlan-delete	`--vlan-range`	For deleting VLAN or VLAN-range
vlan-update	`--vlan-range`	Updates entire VLAN range
num-vrf-update	`--vrf-count`	Mandatory
enable-bd-update	`--enable-bd`	Mandatory
mirror-destination-port-add / mirror-destination-port-delete	`--mirror-destination-port`	Mandatory

PO (Port-Channel) Update Operations

Operation	Flags	Notes
port-add / port-delete	`--port`	Mandatory
lacp-timeout	`--lacp-timeout`	Mandatory
description	`--description`	Mandatory
min-link-count	`--min-link-count`	Mandatory
mtu-add	`--mtu`	Mandatory
mtu-delete	*(none)*	Removes existing custom MTU

VRF Update Operations

Operation	Flags	Notes
local-asn-add	`--local-asn`	Mandatory
local-asn-delete	*(none)*	Removes existing ASN
static-route-bfd-add / static-route-bfd-delete	`--ipv4-static-route-bfd, --ipv6-static-route-bfd`	Either or both can be provided
static-route-add / static-route-delete	`--ipv4-static-route-next-hop, --ipv6-static-route-next-hop`	Either or both can be provided
max-path-add	`--max-path`	Mandatory
max-path-delete	*(none)*	Removes max-path
centralized-router-add	`--centralized-router`	Mandatory
centralized-router-delete	`--centralized-router`	Mandatory
redistribute-add	`--redistribute`	Mandatory
redistribute-delete	`--redistribute`	Mandatory
rh-max-path-add	`--rh-max-path`	Mandatory
rh-max-path-delete	*(none)*	Removes RH max-path
rh-ecmp-update	`--rh-ecmp-enable`	Mandatory
graceful-restart-update	`--graceful-restart-enable`	Mandatory
next-hop-recursion-update	`--next-hop-recursion-enable`	Mandatory
default-information-originate-update	`--default-information-originate-enable`	Mandatory
network-add / network-delete	`--ipv4-network, --ipv4-network-backdoor, --ipv4-network-weight, --ipv4-network-route-map, --ipv6-network, --ipv6-network-backdoor, --ipv6-network-weight, --ipv6-network-route-map`	Either or all can be provided

Operation	Flags	Notes
static-network-add / static-network-delete	`--ipv4-static-network, --ipv4-static-network-distance`	Either or both can be provided
aggregate-address-add / aggregate-address-delete	`--ipv4-aggregate-address, --ipv4-aggregate-as-set, --ipv4-aggregate-summary-only, --ipv4-aggregate-advertise-map, --ipv4-aggregate-suppress-map, --ipv6-aggregate-address, --ipv6-aggregate-as-set, --ipv6-aggregate-summary-only, --ipv6-aggregate-advertise-map, --ipv6-aggregate-suppress-map`	Either or all can be provided

EPG Update Operations

Operation	Flags	Notes
port-group-add	`--port, --po`	Either or both can be provided
	`--switchport-mode`	Mandatory (when the intended EPG doesn't have any existing Port/PO)
	`--switchport-native-vlan-tagging` `--single-homed-bfd-session-type`	Optional flags (when the intended EPG doesn't have any existing Port/PO)
port-group-delete	`--port, --po`	Either one or all can be provided
port-property-add / port-property-delete / port-property-update	`--pp-mac-acl-in, --pp-mac-acl-out, --pp-ip-acl-in, --pp-ip-acl-out, --pp-ipv6-acl-in`	Either one or all can be provided

Operation	Flags	Notes
ctag-range-add	`--ctag-range`	Mandatory
	`--l2-vni, --ctag-description, --bridge-domain, --switchport-native-vlan`	Optional flags in case of L2 EPG
	`--l2-vni, --ctag-description, --bridge-domain, --switchport-native-vlan, --anycast-ip, --anycast-ipv6, --local-ip, --local-ipv6, --ip-icmp-redirect, --ipv6-icmp-redirect, --ip-mtu, --ipv6-nd-mtu, --ipv6-nd-managed-config, --ipv6-nd-other-config, --ipv6-nd-prefix, --ipv6-nd-prefix-valid-lifetime, --ipv6-nd-prefix-preferred-lifetime, --ipv6-nd-prefix-no-advertise, --ipv6-nd-prefix-config-type, --suppress-arp, --suppress-nd, --dhcpv4-relay-address-ip, --dhcpv6-relay-address-ip, --dhcpv4-relay-address-ip-vrf, --dhcpv6-relay-address-ip-vrf, --dhcpv6-relay-address-ip-interface, --dhcpv4-relay-gateway-ip, --dhcpv4-relay-gateway-ip-interface, --dhcpv4-relay-gateway-interface, --dhcpv6-relay-gateway-interface, --dhcpv6-relay-gateway-interface-ip`	Optional flags in case of L3 EPG
ctag-range-delete	`--ctag-range`	Mandatory
network-property-add / network-property-delete / network-property-update	`--np-mac-acl-in, --np-mac-acl-out, --np-ip-acl-in, --np-ip-acl-out, --np-ipv6-acl-in`	Either or all can be provided

Operation	Flags	Notes
vrf-add	`--vrf`	Mandatory
	`--anycast-ip, --anycast-ipv6`	Mandatory for extension typed EPG
	`--local-ip, --local-ipv6`	Mandatory for extension typed EPG
	`--l3-vni, --ip-icmp-redirect, --ipv6-icmp-redirect, --ip-mtu, --ipv6-nd-mtu, --ipv6-nd-managed-config, --ipv6-nd-other-config, --ipv6-nd-prefix, --ipv6-nd-prefix-valid-lifetime, --ipv6-nd-prefix-preferred-lifetime, --ipv6-nd-prefix-no-advertise, --ipv6-nd-prefix-config-type, --suppress-arp, --suppress-nd, --dhcpv4-relay-address-ip, --dhcpv6-relay-address-ip, --dhcpv4-relay-address-ip-vrf, --dhcpv6-relay-address-ip-vrf, --dhcpv6-relay-address-ip-interface, --dhcpv4-relay-gateway-ip, --dhcpv4-relay-gateway-ip-interface, --dhcpv4-relay-gateway-interface, --dhcpv6-relay-gateway-interface, --dhcpv6-relay-gateway-interface-ip`	Optional flags
vrf-delete	*(none)*	Removes existing VRF
local-ip-add / local-ip-delete	`--local-ip, --local-ipv6`	Either or both can be provided
anycast-ip-add / anycast-ip-delete	`--anycast-ip, --anycast-ipv6`	Either or both can be provided
dhcp-relay-address-ip-add / dhcp-relay-address-ip-delete	`--dhcpv4-relay-address-ip, --dhcpv6-relay-address-ip, --dhcpv4-relay-address-ip-vrf, --dhcpv6-relay-address-ip-vrf, --dhcpv6-relay-address-ip-interface`	Either or all can be provided
dhcp-relay-gateway-ip-add / dhcp-relay-gateway-ip-delete	`--dhcpv4-relay-gateway-ip, --dhcpv4-relay-gateway-ip-interface, --dhcpv4-relay-gateway-interface, --dhcpv6-relay-gateway-interface, --dhcpv6-relay-gateway-interface-ip`	Either or all can be provided

Peer-Group Update Operations

Operation	Flags	Notes
desc-update	`--description`	Mandatory
peer-group-add	`--pg-name, --pg-asn`	Mandatory
	All flags used during `peer-group create`	Optional
peer-group-delete	`--pg-name`	Mandatory

Peer Update Operations

Operation	Flags	Notes
desc-update	`--description`	Mandatory
peer-add	`--ipv4-uc-nbr, --ipv6-uc-nbr`	Either or both can be provided for Static Peer
	All flags used during `peer create`	Optional For Static Peer
	`--ipv4-uc-dyn-nbr, --ipv6-uc-dyn-nbr`	Either or both can be provided for the Dynamic Peer
peer-delete	`--ipv4-uc-nbr, --ipv6-uc-nbr`	Either or both can be provided for Static Peer
	`--ipv4-uc-dyn-nbr, --ipv6-uc-dyn-nbr`	Either or both can be provided for the Dynamic Peer

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 endpoints and VRFs, and can also create EPGs of type 'L3-handoff' or 'port-profile', but not 'extension'.

Endpoints are locations where data flows into or out of a network. They consist of hardware components that communicate using protocols like TCP/IP. For example, endpoints may be individual ports or port-channels. Ports are permanent, whereas you can configure and add port-channels when required.

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)

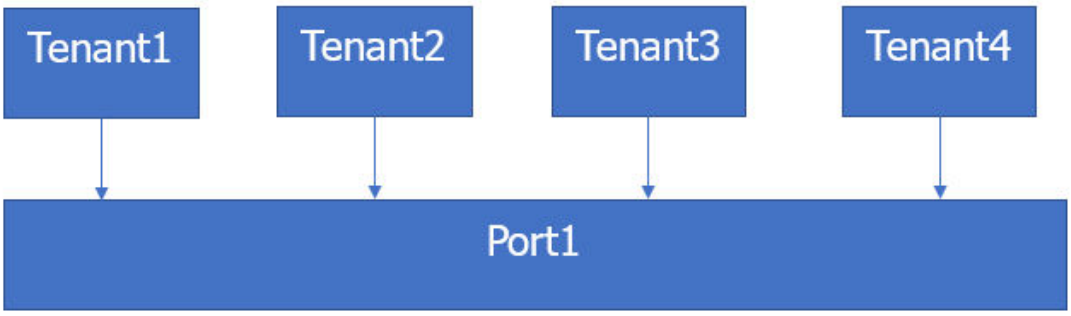

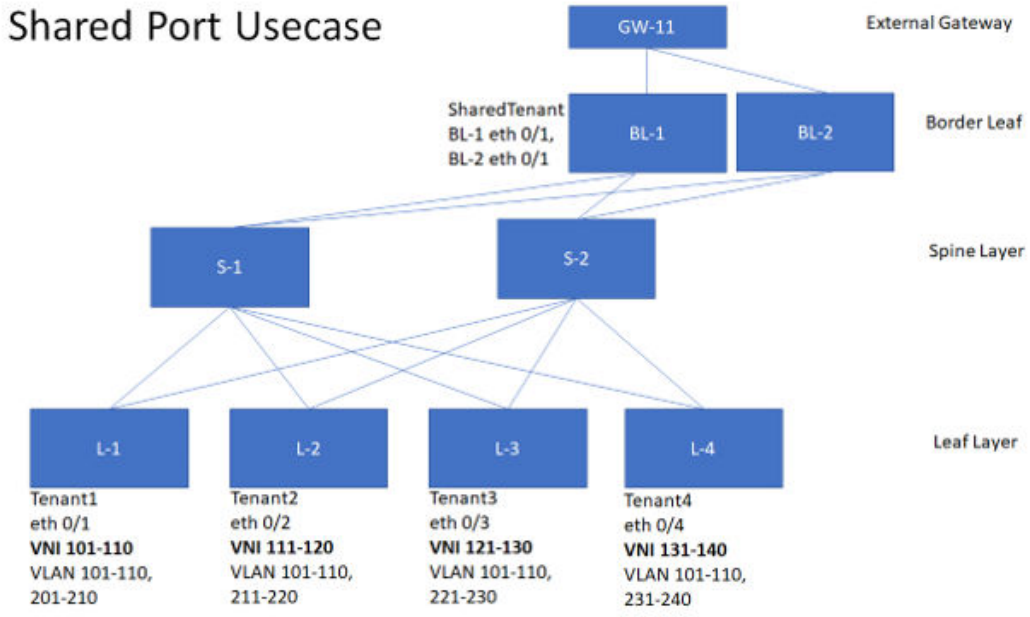

The following examples show the commands and syntax used to configure the Shared Port.

```
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)

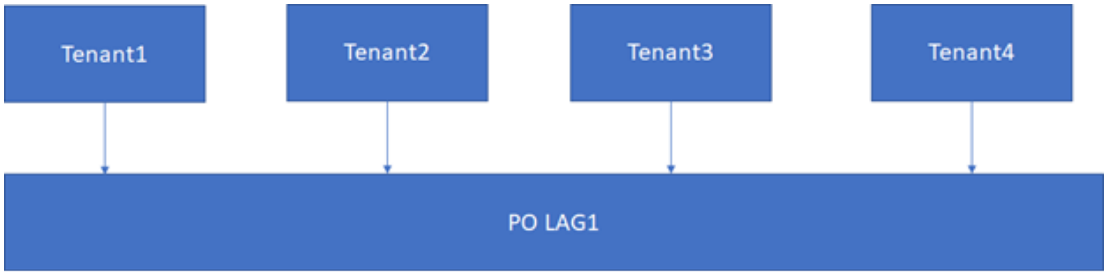

The following examples show the commands and syntax used to configure the Shared Endpoint.

```
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)

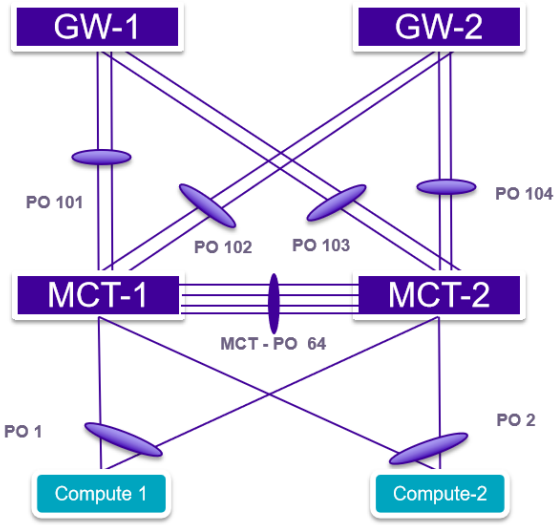

3

Figure 24: 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.

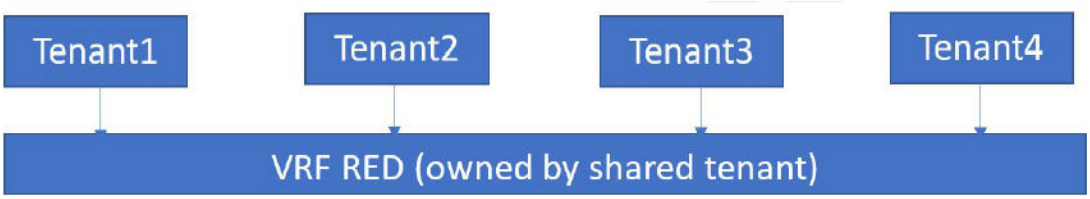

Configure Shared Tenant, Shared VRF, and Private EPG using Shared VRF

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.

About This Task
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 <username> --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
```

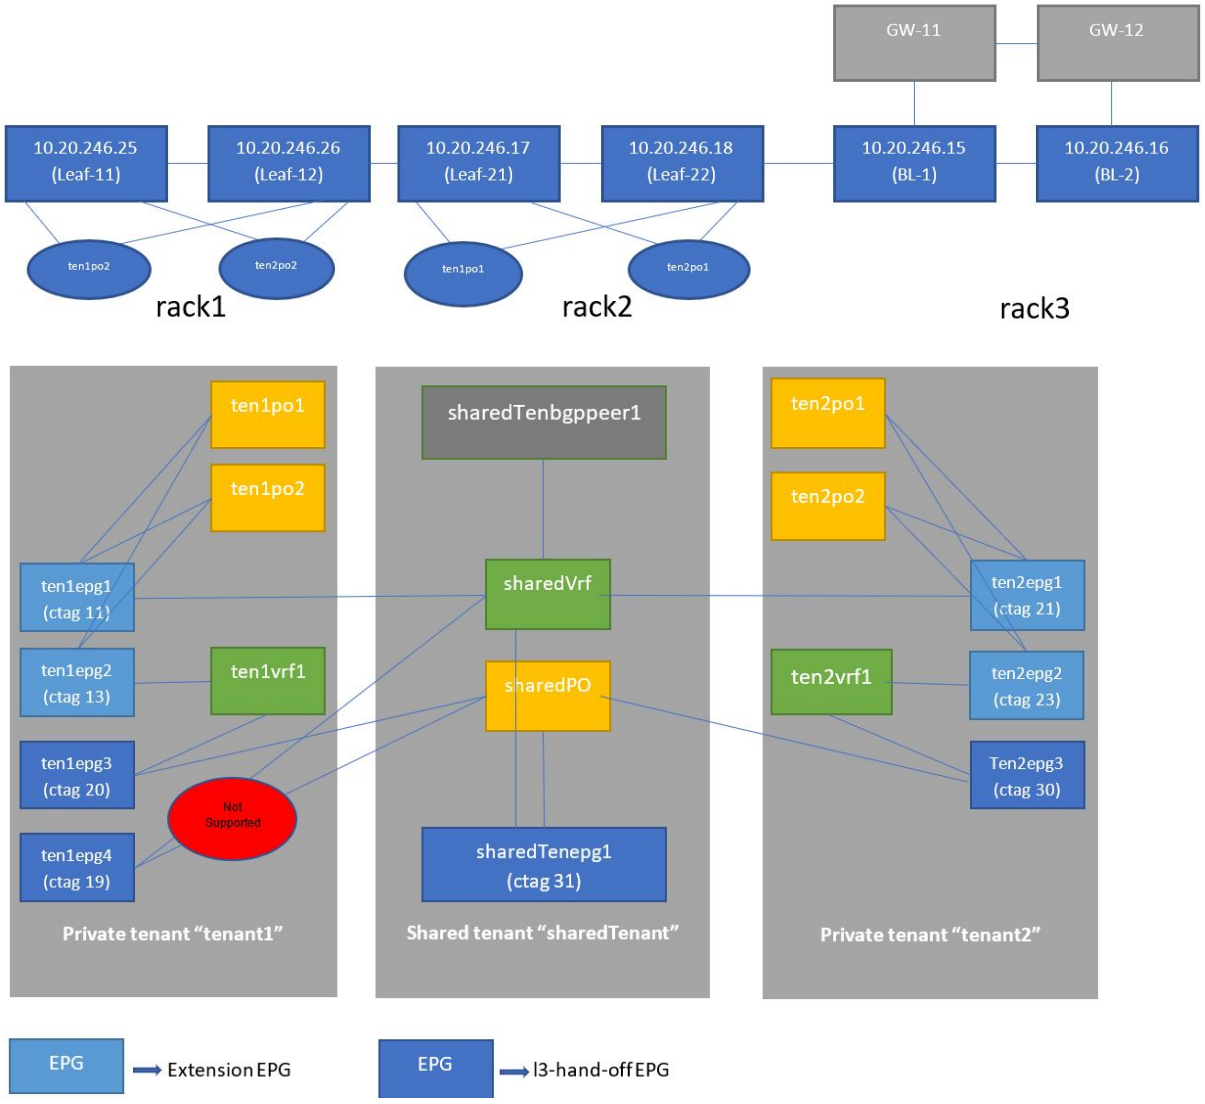

Figure 25: Shared VRF Configuration Overview

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
```

Configure MPLS L3 VPN Inter AS

You can configure the Inter AS.

About This Task

You can configure the Inter AS between SLX 9740 and PE, focusing on:

- LDP for LSP label exchange
- MP-eBGP for VPNv4 and VPNv6 prefix exchanges
- MPLS data plane with a two-label stack for L3 VPN
- MPLS EXP(TC) bits classification and actions
- Authenticated MP-BGP peering with external PEs

Supported Devices: SLX 9740 and SLX 8820

Follow this procedure to configure Inter AS between SLX 9740 and PE.

Procedure

1. Enable or disable switch process restart for inventory services.
   SLX-OS has the following default settings for process restart:
   - MPLS process auto-restart is enabled by default.
   - BGP, OSPFv4, and OSPFv6 process auto-restart are disabled by default.

   To accommodate different deployment needs, you can configure process restart settings for specific processes using the **`efa inventory device setting update`** command. This allows for more control over process restarts. It is an independent command that is pushed to devices on execution of the **`efa inventory device setting update`** command.

```
efa inventory device list
+--------------+-----------+-------+--------------+-------------------------------
+------------+------------+--------+-------------------------+-------------------------
+----------------+-------------+------------------+---------------------------
+-------------------+-----------------+---------+----------+
|  IP Address  | Host Name | Model | Chassis Name |
Firmware            |    ASN     |    Role    | Fabric |    RegistrationTime
|    LastDiscoveryTime    | Cert/Key Saved | Admin State | Maintenance Mode |
Maintenance Mode on Reboot | Syslog Registered | SNMP Registered |  Type  | Location |
+--------------+-----------+-------+--------------+-------------------------------
+------------+------------+--------+-------------------------+-------------------------
+----------------+-------------+------------------+---------------------------
+-------------------+-----------------+---------+----------+
| 10.64.196.61 | SLX       | 4000  | BR-SLX9540   |
20.7.2slxos20.7.2_250505_0600 | 4200065535 | BorderLeaf | fabs   | 2025-08-08 06:38:44
EDT | 2025-08-08 06:47:01 EDT | N              | up          | Disable          |
Disable                     | N               | N               | FABRIC | default  |
+--------------+-----------+-------+--------------+-------------------------------
+------------+------------+--------+-------------------------+-------------------------
+----------------+-------------+------------------+---------------------------
+-------------------+-----------------+---------+----------+
| 10.64.196.62 | SLX62     | 4000  | BR-SLX9540   |
20.7.2slxos20.7.2_250505_0600 | 4200065535 | BorderLeaf | fabs   | 2025-08-08 06:38:44
EDT | 2025-08-08 06:47:01 EDT | N              | up          | Disable          |
Disable                     | N               | N               | FABRIC | default  |
+--------------+-----------+-------+--------------+-------------------------------
+------------+------------+--------+-------------------------+-------------------------
+----------------+-------------+------------------+---------------------------
+-------------------+-----------------+---------+----------+
Device Details
```

   For information on command syntax and parameter description, see *ExtremeCloud Orchestrator Command Reference Guide*.

   a. Disable the MPLS process restart.
      Before updating the process restart setting, disable MPLS process restart on devices using the following command:

```
efa inventory device setting update --ip 10.64.196.61,10.64.196.62 --process-
restart MPLS,NO
+--------------+----------------------+---------+-------+-------+
|  IP ADDRESS  |         NAME         | STATUS  | VALUE | ERROR |
+--------------+----------------------+---------+-------+-------+
| 10.64.196.61 | Process Restart MPLS | Success | NO    |       |
```

```
+--------------+----------------------+---------+-------+-------+
| 10.64.196.62 | Process Restart MPLS | Success | NO    |       |
+--------------+----------------------+---------+-------+-------+
```

- Verify the MPLS process restart status on devices.

```
efa inventory device setting show --ip 10.64.196.61 | grep -I "process restart"
+---------------------------------+-------+-------------+
|              NAME               | VALUE |  APP STATE  |
+---------------------------------+-------+-------------+
| Process Restart Enabled On      | No    | cfg-in-sync |
| MPLS                            |       |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On BGP | No    |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On      | No    |             |
| ISIS                            |       |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On      | No    |             |
| OSPFv2                          |       |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On      | No    |             |
| OSPFv3                          |       |             |
+---------------------------------+-------+-------------+
```

- Verify the configuration on SLX devices:

```
show running-config ha
ha
 no process-restart mpls         ----- MPLS process restart is disabled.
 no process-restart bgp
 no process-restart isis
 no process-restart ospfv2
 no process-restart ospfv3
!
SLX#
```

b. Enable the MPLS process restart.

Before updating the process restart setting, disable MPLS process restart on devices using the following command:

```
efa inventory device setting update  --ip 10.64.196.61,10.64.196.62 --process-
restart MPLS,YES
+--------------+----------------------+---------+-------+-------+
|  IP ADDRESS  |         NAME         | STATUS  | VALUE | ERROR |
+--------------+----------------------+---------+-------+-------+
| 10.64.196.61 | Process Restart MPLS | Success | Yes   |       |
+--------------+----------------------+---------+-------+-------+
| 10.64.196.62 | Process Restart MPLS | Success | Yes   |       |
+--------------+----------------------+---------+-------+-------+
```

- Verify the MPLS process restart status on devices.

```
efa inventory device setting show --ip 10.64.196.61 | grep -i "process restart"
+---------------------------------+-------+-------------+
|              NAME               | VALUE |  APP STATE  |
+---------------------------------+-------+-------------+
...
| Process Restart Enabled On      | Yes   | cfg-in-sync |
| MPLS                            |       |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On BGP | No    |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On      | No    |             |
| ISIS                            |       |             |
+---------------------------------+-------+-------------+
| Process Restart Enabled On      | No    |             |
| OSPFv2                          |       |             |
```

```
+------------------------------+-------+-------------+
| Process Restart Enabled On     | No    |             |
| OSPFv3                         |       |             |
+------------------------------+-------+-------------+
--- Time Elapsed: 68.882949ms ---
```

- Verify the configuration on SLX devices:

```
show running-config ha
ha
 process-restart mpls         -----  process-restart enabled.
 no process-restart bgp
 no process-restart isis
 no process-restart ospfv2
 no process-restart ospfv3
!
SLX#
```

2. Configure MPLS on a fabric.

   You can enable or disable MPLS at the fabric level and configure the MPLS loopback ID and IP pool. The setting determines whether the MPLS L3VPN-based L3 handoff can be used. Additionally, you can configure the loopback ID and IP pool for the router MPLS ID, providing flexibility in your MPLS setup. The default loopback scheme is Uniform.

   a. MPLS Disabled to Enabled on a newly created fabric:

      On a newly created fabric, run the following command to create fabric, enable MPLS, and configure loopback settings:

```
efa fabric create -name "fab2" -type clos --stage 3

efa fabric setting show --name fab2
+----------------------------+------------------------+
|            NAME            |         VALUE          |
+----------------------------+------------------------+
| Fabric Name                | fab2                   |
+----------------------------+------------------------+
| Loopback Scheme            | granular               |
+----------------------------+------------------------+
| Loopback IP Range          | 172.31.254.0/24        |
+----------------------------+------------------------+
| RouterID Loopback IP Range | 172.31.128.0/24        |
+----------------------------+------------------------+
| VTEP Loopback IP Range     | 172.31.192.0/24        |
+----------------------------+------------------------+
| Loopback Port Number       | 1                      |
+----------------------------+------------------------+
| VTEP Loopback Port Number  | 3                      |
+----------------------------+------------------------+
| MPLS Router Enable         | No                     |
+----------------------------+------------------------+

efa fabric setting update  --mpls-router-enable yes --name fab2
Error :  VTEPLoopBackIPRange 172.31.192.0/24 Overlaps with
MPLSRouterIDLoopBackIPRange 172.31.192.0/24

efa fabric delete --name fab2 -force

efa fabric create --name fab2 --type non-clos

efa fabric setting show --name fab2
+----------------------------+------------------------+
|            NAME            |         VALUE          |
+----------------------------+------------------------+
| Fabric Name                | fab2                   |
+----------------------------+------------------------+
```

```
| Loopback Scheme            | uniform               |
+----------------------------+-----------------------+
| Loopback IP Range          | 172.31.254.0/24       |
+----------------------------+-----------------------+
| RouterID Loopback IP Range | 172.31.128.0/24       |
+----------------------------+-----------------------+
| VTEP Loopback IP Range     | 172.31.64.0/24        |
+----------------------------+-----------------------+
| Loopback Port Number       | 1                     |
+----------------------------+-----------------------+
| VTEP Loopback Port Number  | 2                     |
+----------------------------+-----------------------+
| MPLS Router Enable         | No                    |
+----------------------------+-----------------------+

           |
+---------------------------------+-----------------------+
efa fabric setting update --loopback-scheme granular --name fab2
efa fabric setting update  --mpls-router-enable yes --name fab2
           --mpls-router-id-loopback-ip-range 172.31.198.0/23 --mpls-loopback-port-
number 7

efa fabric setting show --name fab2
+---------------------------------+-----------------------+
|              NAME               |         VALUE         |
+---------------------------------+-----------------------+
| Fabric Name                     | fab2                  |
+---------------------------------+-----------------------+
| Loopback Scheme                 | granular              |
+---------------------------------+-----------------------+
| MPLS Router ID Loopback IP      | 172.31.198.0/23       |
| range                           |                       |
+---------------------------------+-----------------------+
| MPLS Loopback Port Number       | 7                     |
+---------------------------------+-----------------------+
| MPLS Router Enable              | Yes                   |
+---------------------------------+-----------------------+

efa fabric device add-bulk --name fab2 --rack room1-rack1
           --border-leaf-ip 10.64.160.11,10.64.160.38 --border-leaf-rack rack1 --
username admin --password password

efa fabric configure --name fab2

efa fabric show

Fabric Name: default, Fabric Description: Default Fabric, 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 | MPLS LB ID |
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+------------+
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+------------+

Fabric Name: fab2, Fabric Description: , Fabric
Type: non-clos, Fabric Status: created, Fabric Health: Red
+--------------+-------+-----------+------------+------------+-----------------
+-----------+-------------------+----------------------------------------+---------
+-------+------------+
|  IP ADDRESS  | RACK  | HOST NAME
|    ASN     |    ROLE    |  DEVICE STATE   | APP STATE | CONFIG
```

```
GEN REASON |           PENDING CONFIGS            | VTLB ID | LB ID | MPLS LB ID |
+--------------+-------+-----------+------------+------------+-----------------
+-----------+------------------+---------------------------------------+---------
+-------+------------+
| 10.64.160.11 | rack1 | SLX       | 4200065535
| BorderLeaf | not provisioned | cfg ready | DA
| SYSP-C,MCT-C,MCT-PA,BGP-C,EVPN-C,O-C | 2       | 1     | 7          |
| 10.64.160.38 | rack1 | SLX       | 4200065535
| BorderLeaf | not provisioned | cfg ready | DA
| SYSP-C,MCT-C,MCT-PA,BGP-C,EVPN-C,O-C | 2       | 1     | 7          |
+--------------+-------+-----------+------------+------------+-----------------
+-----------+------------------+---------------------------------------+---------
+-------+------------+



efa tenant create --name "t2" --type shared
--vlan-range 2-4000 --vrf-count 100 --port 10.64.160.38[0/3-5],10.64.160.11[0/3-5]

efa tenant vrf create --tenant "t2" --name default-vrf --routing-type "centralized"


efa fabric setting update --mpls-router-enable no --name fab2
Error :  MPLS Router cannot be disabled because
the fabric has some tenants with default VRF configurations.
```

- Verify the following configurations on SLX devices:

BL1

```
interface Loopback 1
 ip address 172.31.128.97/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.64/32
 no shutdown
!
interface Loopback 7
 ip address 172.31.198.247/32
 no shutdown
!
router bgp
 local-as 66000
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
!
```

BL2

```
interface Loopback 1
 ip address 172.31.128.204/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.64/32
 no shutdown
!
interface Loopback 7
 ip address 172.31.198.177/32
 no shutdown
!
router bgp
 local-as 66000
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
!
```

b. MPLS Enabled to Disabled on a newly created fabric:

On a newly created fabric, run the following command to delete tenant VRFs, disable MPLS, and verify fabric settings:

```
efa tenant vrf delete --name default-vrf --tenant t2

efa fabric setting update --mpls-router-enable no --name fab2

efa fabric show

Fabric Name: default, Fabric Description: Default Fabric, 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 | MPLS LB ID |
+------------+-----+-----------+-----+------+--------------+-----------
+------------------+----------------+---------+-------+------------+
+------------+-----+-----------+-----+------+--------------+-----------
+------------------+----------------+---------+-------+------------+

Fabric Name: fab2, Fabric Description: , Fabric
Type: non-clos, Fabric Status: settings-updated, Fabric Health: Red

Updated Fabric Settings: MPLS-ENABLE

+--------------+-------+-----------+------------+------------+--------------
+-------------+------------------+-----------------+---------+-------+------------+
|  IP ADDRESS  | RACK  | HOST NAME |    ASN     |    ROLE    | DEVICE STATE
|  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID | MPLS LB ID |
+--------------+-------+-----------+------------+------------+--------------
+-------------+------------------+-----------------+---------+-------+------------+
| 10.64.160.11 | rack1 | SLX       | 4200065535 | BorderLeaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     | 7          |
| 10.64.160.38 | rack1 | SLX       | 4200065535 | BorderLeaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     | 7          |
+-------------+-------+-----------+------------+------------+--------------
+-------------+------------------+-----------------+---------+-------+------------+
FABRIC SETTING:


efa fabric setting show --name fab2
+----------------------------+-----------------------+
|            NAME            |         VALUE         |
+----------------------------+-----------------------+
| Fabric Name                | fab2                  |
+----------------------------+-----------------------+
| Link IP Range              | 10.10.10.0/23         |
+----------------------------+-----------------------+
| Loopback Scheme            | granular              |
+----------------------------+-----------------------+
| Loopback IP Range          | 172.31.254.0/24       |
+----------------------------+-----------------------+
| RouterID Loopback IP Range | 172.31.128.0/24       |
+----------------------------+-----------------------+
| VTEP Loopback IP Range     | 172.31.64.0/24        |
+----------------------------+-----------------------+
| Loopback Port Number       | 1                     |
+----------------------------+-----------------------+
| VTEP Loopback Port Number  | 2                     |
+----------------------------+-----------------------+
| Rack ASN Block             | 4200000000-4200065534 |
+----------------------------+-----------------------+
| Rack Border Leaf ASN Block | 4200065535-4200065635 |
+----------------------------+-----------------------+
| Single Rack Deployment     | No                    |
+----------------------------+-----------------------+
| P2P IP Type                | numbered              |
+----------------------------+-----------------------+
| MPLS Router Enable         | No                    |
+----------------------------+-----------------------+

efa fabric configure --name fab2
```

- Verify the following configurations on SLX devices:

BL1

```
interface Loopback 1
 ip address 172.31.128.97/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.64/32
 no shutdown
!
router bgp
 local-as 66000
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
!
```

BL2

```
interface Loopback 1
 ip address 172.31.128.204/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.64/32
 no shutdown
!
router bgp
 local-as 66000
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
!
```

c. MPLS Disabled to Enabled on an existing provisioned fabric:

On an existing provisioned fabric, run the following command to enable MPLS, configure loopback settings, and verify fabric settings:

```
efa fabric setting update  --mpls-router-enable yes --name fab2
           --mpls-router-id-loopback-ip-range 172.31.172.0/24 --mpls-loopback-port-
number 9


efa fabric setting show --name fab2
+--------------------------------+------------------------+
|              NAME              |         VALUE          |
+--------------------------------+------------------------+
| Fabric Name                    | fab2                   |
+--------------------------------+------------------------+
| Link IP Range                  | 10.10.10.0/23          |
+--------------------------------+------------------------+
| Loopback Scheme                | granular               |
+--------------------------------+------------------------+
| Loopback IP Range              | 172.31.254.0/24        |
+--------------------------------+------------------------+
| RouterID Loopback IP Range     | 172.31.128.0/24        |
+--------------------------------+------------------------+
| VTEP Loopback IP Range         | 172.31.64.0/24         |
+--------------------------------+------------------------+
| MPLS Router ID Loopback IP     | 172.31.172.0/24        |
| range                          |                        |
+--------------------------------+------------------------+
| Loopback Port Number           | 1                      |
+--------------------------------+------------------------+
| VTEP Loopback Port Number      | 2                      |
+--------------------------------+------------------------+
| MPLS Loopback Port Number      | 9                      |
+--------------------------------+------------------------+
| Rack ASN Block                 | 4200000000-4200065534 |
+--------------------------------+------------------------+
| Rack Border Leaf ASN Block     | 4200065535-4200065635 |
+--------------------------------+------------------------+
| Single Rack Deployment         | No                     |
+--------------------------------+------------------------+
| P2P IP Type                    | numbered               |
+--------------------------------+------------------------+
| MPLS Router Enable             | Yes                    |
+--------------------------------+------------------------+
```

```
fa fabric show

Fabric Name: default, Fabric Description: Default Fabric, 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 | MPLS LB ID |
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+------------+
+------------+-----+-----------+-----+------+--------------+-----------
+-------------------+-----------------+---------+-------+------------+

Fabric Name: fab2, Fabric Description: , Fabric
Type: non-clos, Fabric Status: settings-updated, Fabric Health: Red

Updated Fabric Settings: MPLS-ENABLE

+--------------+-------+-----------+------------+------------+--------------
+--------------+-------------------+-----------------+---------+-------+------------+
|  IP ADDRESS  | RACK  | HOST NAME |    ASN     |    ROLE    | DEVICE STATE
|  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID | MPLS LB ID |
+--------------+-------+-----------+------------+------------+--------------
+--------------+-------------------+-----------------+---------+-------+------------+
| 10.64.160.11 | rack1 | SLX       | 4200065535 | BorderLeaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |            |
| 10.64.160.38 | rack1 | SLX       | 4200065535 | BorderLeaf | provisioned
| cfg in-sync | NA                | NA              | 2       | 1     |            |
+--------------+-------+-----------+------------+------------+--------------
+--------------+-------------------+-----------------+---------+-------+------------+

efa fabric configure --name fab2
```

- Verify the following configurations on SLX devices:

BL1

```
interface Loopback 1
 ip address 172.31.128.97/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.64/32
 no shutdown
!
interface Loopback 9
 ip address 172.31.172.87/32
 no shutdown
!
router bgp
 local-as 66000
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
!
```

BL2

```
interface Loopback 1
 ip address 172.31.128.204/32
 no shutdown
!
interface Loopback 2
 ip address 172.31.64.64/32
 no shutdown
!
interface Loopback 9
 ip address 172.31.172.87/32
 no shutdown
!
router bgp
 local-as 66000
 address-family ipv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
!
```

3. Configure a shared tenant.

   To utilize shared resources such as VRFs, EPGs, and BGP peers, configure a shared tenant. This concept is already supported in the XCO tenant, allowing for efficient management of shared resources under a single entity.

```
efa tenant create --name "SharedTen1" --type shared
                  --vlan-range 4000-4001 --vrf-count 1 --port
10.64.164.44[0/1-4],10.64.164.45[0/1-4]
```

4. Configure shared Port-Channel.

   To configure a shared port-channel for carrying L3 Handoff traffic between the BL and GW pair, run the following command:

```
efa tenant po create --name SharedTen1Po11 --tenant SharedTen1 --port
10.64.164.45[0/1-2]
                     --speed 40Gbps --negotiation active -description "BL1_GW1 Member
interface"

efa tenant po create --name SharedTen1Po12 --tenant SharedTen1 --port
10.64.164.45[0/3-4]
                     --speed 40Gbps --negotiation active -description "BL1_GW2 Member
interface"
efa tenant po create --name SharedTen1Po21 --tenant SharedTen1 --port
10.64.164.44[0/1-2]
                     --speed 40Gbps --negotiation active -description "BL2_GW1 Member
interface"

efa tenant po create --name SharedTen1Po22 --tenant SharedTen1 --port
10.64.164.44[0/3-4]
```

```
                    --speed 40Gbps --negotiation active -description "BL2_GW2 Member
interface"
```

- Verify the following configurations on SLX devices:

BL1

```
vrf Ten1Vrf1
 rd 172.31.254.190:1
 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
 !
!
vrf Ten2Vrf1
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
 !
 address-family ipv6 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
 !
!
vrf Ten1Vrf2
 rd 172.31.254.190:1
 address-family ipv4 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
!
vrf Ten2Vrf2
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
!
```

BL2

```
vrf Ten1Vrf1
 rd 172.31.254.190:1
 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
 !
!
vrf Ten2Vrf1
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
 !
 address-family ipv6 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
 !
!
vrf Ten1Vrf2
 rd 172.31.254.190:1
 address-family ipv4 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
!
vrf Ten2Vrf2
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
!
```

5. Configure shared VRF under the shared tenant with static routes.

   To configure a shared VRF for carrying L3 Handoff traffic between the BL and GW pair, with static routes, run the following command:

```
efa tenant vrf create --name default-vrf --tenant SharedTen1 --routing-type centralized

                      --ipv4-static-route-next-hop
10.64.164.45,12.12.12.12/32,99.99.99.2
                      --ipv4-static-route-next-hop
10.64.164.45,14.14.14.14/32,88.88.88.2
                      --ipv4-static-route-next-hop
10.64.164.44,13.13.13.13/32,77.77.77.2
                      --ipv4-static-route-next-hop
10.64.164.44,15.15.15.15/32,66.66.66.2
```

6. Configure shared L3 Handoff EPG.

   Configure a shared L3 Handoff EPG for configuring Local IP address needed for the BGP peering between the BL and GW pair.

   a. To create an EPG, run the following command:

```
efa tenant epg create --name SharedTen1L3HandOffEPG11 --type l3-hand-off
                      --vrf default-vrf
                      --switchport-mode trunk --po SharedTen1PO11
                      --ctag-range 4000 --local-ip 4000,10.64.164.45:99.99.99.1/24
efa tenant epg create --name SharedTen1L3HandOffEPG12 --type l3-hand-off
                      --vrf default-vrf
                      --switchport-mode trunk --po SharedTen1PO12
                      --ctag-range 4001 --local-ip 4000,10.64.164.45:88.88.88.1/24
efa tenant epg create --name SharedTen1L3HandOffEPG21 --type l3-hand-off
                      --vrf default-vrf
                      --switchport-mode trunk --po SharedTen1PO21
                      --ctag-range 4000 --local-ip 4000,10.64.164.44:77.77.77.1/24
efa tenant epg create --name SharedTen1L3HandOffEPG22 --type l3-hand-off
                      --vrf default-vrf
                      --switchport-mode trunk --po SharedTen1PO22
                      --ctag-range 4001 --local-ip 4000,10.64.164.44:66.66.66.1/24
```

   b. Verify the following configurations on SLX devices:

BL1

```
router mpls
 ldp
  lsr-id 11.11.11.11
 !
 mpls-interface ve 4000
  ldp-enable
 !
 mpls-interface ve 4001
  ldp-enable
 !
!
```

BL2

```
router mpls
 ldp
  lsr-id 11.11.11.12
 !
 mpls-interface ve 4000
  ldp-enable
 !
 mpls-interface ve 4001
  ldp-enable
 !
!
```

BL1

```
interface Port-channel 1
 speed 40000
 cluster-track
 description BL1_GW1 Member interface
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 4000
 switchport trunk tag native-vlan
 no shutdown
```

BL2

```
interface Port-channel 1
 speed 40000
 cluster-track
 description BL2_GW1 Member interface
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 4000
 switchport trunk tag native-vlan
 no shutdown
```

```
!
interface Port-channel 2
 speed 40000
 cluster-track
 description BL1_GW2 Member interface
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 4001
 switchport trunk tag native-vlan
 no shutdown
!
vlan 4000
 router-interface Ve 4000
!
vlan 4001
 router-interface Ve 4001
!
interface Ve 4000
 ip address 99.99.99.1/24
 no shutdown
!
interface Ve 4001
 ip address 88.88.88.1/24
 no shutdown
!
ip route 12.12.12.12/32 99.99.99.2
ip route 14.14.14.14/32 88.88.88.2
```

```
!
interface Port-channel 2
 speed 40000
 cluster-track
 description BL2_GW2 Member interface
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 4001
 switchport trunk tag native-vlan
 no shutdown
!
vlan 4000
 router-interface Ve 4000
!
vlan 4001
 router-interface Ve 4001
!
interface Ve 4000
 ip address 77.77.77.1/24
 no shutdown
!
interface Ve 4001
 ip address 66.66.66.1/24
 no shutdown
!
ip route 13.13.13.13/32 77.77.77.2
ip route 15.15.15.15/32 66.66.66.2
```

c. Add a Device to XCO.

Use the **port-group-add** command with local-ip (required for router MPLS). Adding ports to existing devices is optional.

```
efa tenant epg update --name SharedTen1L3HandOffEPG11 -tenant SharedT1 --operation
port-group-add -port 10.64.164.45[0/28] --local-ip 4000,10.64.164.45:99.99.99.1/24
```

d. Validations & Errors

Command	Error
Missing local-ip: local-ip is mandatory for router MPLS to be pushed. `efa tenant epg update --name SharedTen1L3HandOffEPG11 -tenant SharedT1 --operation port-group-add -port 10.64.164.45[0/28]`	local-ip is mandatory for router MPLS enabled
EPG without network properties: If EPG is created without network properties, device-specific network properties cannot be given. `efa tenant epg update --name SharedTen1L3HandOffEPG11 -tenant SharedT1 --operation port-group-add -port 10.64.164.45[0/28] --local-ip 4000,10.64.164.45:99.99.99.1/24`	PortGroup of EPG SharedTen1L3HandOffEPG11 is configured without network properties. Hence device specific network properties cannot be given
Mismatched ctag: Ctag range must match the EPG configuration `efa tenant epg update --name SharedTen1L3HandOffEPG11 -tenant SharedT1 --operation port-group-add -port 10.64.164.45[0/28] --local-ip 5000,10.64.164.45:99.99.99.1/24`	Given ctag is not the range

7. Configure shared BGP peer for peering with external gateway.

a. Run the following command:

```
efa tenant service bgp peer create -h

  --ipv4-uc-nbr-activate stringArray                                IPv4 Unicast
Neighbor Activate in the format device-ip,vrf-name:ipv4-neighbor,activate(true/
false). Only for SLXOS MPLS BGP peers, Not Supported for EXTREME-OS-ONE devices.
  --ipv4-uc-nbr-activate-in-ipvpnv4-uc-af stringArray               IPv4 Unicast
Neighbor in the format device-ip,vrf-name:ipv4-neighbor,activate(true/false). Only
for SLXOS MPLS BGP peers, Not Supported for EXTREME-OS-ONE devices.
  --ipv4-uc-nbr-activate-in-ipvpnv6-uc-af stringArray               IPv4 Unicast
Neighbor in the format device-ip,vrf-name:ipv4-neighbor,activate(true/false). Only
for SLXOS MPLS BGP peers, Not Supported for EXTREME-OS-ONE devices.

efa tenant service bgp peer create --tenant SharedTen1 --name
SharedTen1BgpPeerFromBL1ToGw1AndGw2
               --ipv4-uc-nbr 10.64.164.45,default-vrf: 88.88.88.2,1000
               --ipv4-uc-nbr-bfd 10.64.164.45,default-vrf: 88.88.88.2,true,300,300,3
               --ipv4-uc-nbr-activate-in-ipv4-uc-af 10.64.164.45,default-vrf:
88.88.88.2,false
               --ipv4-uc-nbr-activate-in-ipvpnv4-uc-af 10.64.164.45,default-vrf:
88.88.88.2,true
               --ipv4-uc-nbr-activate-in-ipvpnv6-uc-af 10.64.164.45,default-vrf:
88.88.88.2,true
               --ipv4-uc-nbr 10.64.164.45,default-vrf: 99.99.99.2,1000
```

```
               --ipv4-uc-nbr-bfd 10.64.164.45,default-vrf: 99.99.99.2,true,300,300,3
               --ipv4-uc-nbr-activate-in-ipv4-uc-af 10.64.164.45,default-vrf:
99.99.99.2,false
               --ipv4-uc-nbr-activate-in-ipvpnv4-uc-af 10.64.164.45,default-vrf:
99.99.99.2,true
               --ipv4-uc-nbr-activate-in-ipvpnv6-uc-af 10.64.164.45,default-vrf:
99.99.99.2,true

efa tenant service bgp peer create --tenant SharedTen1 --name
SharedTen1BgpPeerFromBL2ToGw1AndGw2
               --ipv4-uc-nbr 10.64.164.44,default-vrf: 66.66.66.2,1000
               --ipv4-uc-nbr-bfd 10.64.164.44,default-vrf: 66.66.66.2,true,300,300,3
               --ipv4-uc-nbr-activate-in-ipv4-uc-af 10.64.164.44,default-vrf:
66.66.66.2,false
               --ipv4-uc-nbr-activate-in-ipvpnv4-uc-af 10.64.164.44,default-vrf:
66.66.66.2,true
               --ipv4-uc-nbr-activate-in-ipvpnv6-uc-af 10.64.164.44,default-vrf:
66.66.66.2,true

               --ipv4-uc-nbr 10.64.164.44,default-vrf: 77.77.77.2,1000
               --ipv4-uc-nbr-bfd 10.64.164.44,default-vrf: 77.77.77.2,true,300,300,3
               --ipv4-uc-nbr-activate-in-ipv4-uc-af 10.64.164.44,default-vrf:
77.77.77.2,false
               --ipv4-uc-nbr-activate-in-ipvpnv4-uc-af 10.64.164.44,default-vrf:
77.77.77.2,true
               --ipv4-uc-nbr-activate-in-ipvpnv6-uc-af 10.64.164.44,default-vrf:
77.77.77.2,true
```

b. Verify the following configurations on SLX devices:

BL1

```
router bgp
 local-as 66000
 compare-routerid

 neighbor 88.88.88.2 remote-as 1000
 neighbor 88.88.88.2 password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor 88.88.88.2 bfd
 neighbor 88.88.88.2 bfd interval
300 min-rx 300 multiplier 3

 neighbor 99.99.99.2 remote-as 1000
 neighbor 99.99.99.2 password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor 99.99.99.2 bfd
 neighbor 99.99.99.2 bfd interval
300 min-rx 300 multiplier 3

 address-family ipv4 unicast
  no neighbor 88.88.88.2 activate
  no neighbor 99.99.99.2 activate
 !

 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
 address-family vpnv4 unicast
  neighbor 88.88.88.2 activate
  neighbor 88.88.88.2 send-
community extended
  neighbor 99.99.99.2 activate
  neighbor 99.99.99.2 send-
community extended
 !
 address-family vpnv6 unicast
  neighbor 88.88.88.2 activate
  neighbor 88.88.88.2 send-
community extended
  neighbor 99.99.99.2 activate
  neighbor 99.99.99.2 send-
community extended
 !
!
```

BL2

```
router bgp
 local-as 66000
 compare-routerid

 neighbor 66.66.66.2 remote-as 1000
 neighbor 66.66.66.2 password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor 66.66.66.2 bfd
 neighbor 66.66.66.2 bfd interval
300 min-rx 300 multiplier 3

 neighbor 77.77.77.2 remote-as 1000
 neighbor 77.77.77.2 password
$9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor 77.77.77.2 bfd
 neighbor 77.77.77.2 bfd interval
300 min-rx 300 multiplier 3

 address-family ipv4 unicast
  no neighbor 66.66.66.2 activate
  no neighbor 77.77.77.2 activate
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
 !
 address-family vpnv4 unicast
  neighbor 66.66.66.2 activate
  neighbor 66.66.66.2 send-
community extended
  neighbor 77.77.77.2 activate
  neighbor 77.77.77.2 send-
community extended
 !
 address-family vpnv6 unicast
  neighbor 66.66.66.2 activate
  neighbor 66.66.66.2 send-
community extended
  neighbor 77.77.77.2 activate
  neighbor 77.77.77.2 send-
community extended
 !
!
```

8. Configure multiple private tenants.

   You can configure multiple private tenants, each with its own VRFs and EPGs.

The `mpls-l3-handoff` setting controls non-EVPN RTs and BGP next-hop interface loopbacks on devices. By default, this setting is disabled (false). When enabled (true), both configurations are applied.

Note

- The mpls-l3-handoff setting can only be updated for VRFs not yet configured on devices (non-dev VRFs)
- Prerequisite: MPLS must be enabled at the fabric level
- If MPLS L3 handoff is disabled, configuring non-EVPN RTs will result in an error.

a. To configure a tenant, tenant VRF, and tenant EPB, run the following command:

```
efa tenant create --name "Ten1" --type private --vlan-range 10-20 --vrf-count 5 --
port 10.64.196.56[0/1-8],10.64.196.57[0/1-8]
efa tenant create --name "Ten2" --type private --vlan-range 21-30 --vrf-count 5 --
port 10.64.196.56[0/10-18],10.64.196.57[0/10-18]
efa tenant vrf create --tenant "Ten1" --name "Ten1Vrf1" --routing-type "centralized"
          --rt-type import --rt 101:101 --rt-type export --rt 101:101 --max-path 8
--redistribute connected --mpls-l3-handoff false
efa tenant vrf create --tenant "Ten2" --name "Ten2Vrf1" --routing-type " centralized
          --rt-type import --rt 102:102 --rt-type export --rt 102:102 --max-path 8
--redistribute connected
efa tenant vrf create --tenant "Ten1" --name "Ten1Vrf2" --routing-type "centralized"
          --rt-type import --rt 103:103 --rt-type export --rt 103:103 --max-path 8
--redistribute connected --non-evpn-rt 201:201 --non-evpn-rt-type both --mpls-l3-
handoff true
efa tenant vrf create --tenant "Ten2" --name "Ten2Vrf2" --routing-type "centralized"
          --max-path 8 --redistribute connected --non-evpn-rt 202:202 --non-evpn-
rt-type import --non-evpn-rt 202:202 --non-evpn-rt-type export --mpls-l3-handoff
true
efa tenant epg create --tenant "Ten1" --name "Ten1Epg1" --type extension
          --switchport-mode trunk --port 10.64.196.56[0/6] --vrf Ten1Vrf1 --ctag-
range 10 --anycast-ip 10:10.10.10.10/24
efa tenant epg create --tenant "Ten2" --name "Ten2Epg1" --type extension
          --switchport-mode trunk --port 10.64.196.56[0/16] --vrf Ten2Vrf1 --ctag-
range 22 --anycast-ip 22:10.10.10.10/24
efa tenant epg create --tenant "Ten1" --name "Ten1Epg2" --type extension
          --switchport-mode trunk --port 10.64.196.56[0/7] --vrf Ten1Vrf2 --ctag-
range 11 --anycast-ip 11:10.10.10.10/24
efa tenant epg create --tenant "Ten2" --name "Ten2Epg2" --type extension
          --switchport-mode trunk --port 10.64.196.56[0/17] --vrf Ten2Vrf2 --ctag-
range 10 --anycast-ip 23:10.10.10.10/24
efa tenant vrf show --tenant Ten1 --detail
================================================================================
================================================================================
Name                                    : Ten1Vrf1
Tenant                                  : Ten1
Routing Type                            : centralized
Centralized Routers                     : 10.64.196.56
                        : 10.64.196.57
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 MPLS L3 Handoff                  : false
Enable Graceful Restart                 : false
Enable NextHop Recursion                : false
Enable Default Information Originate    : false
EVPN Route Target                       : import 101:101
                                        : export 101:101
Non-EVPN Route Target                   :
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

================================================================================
================================================================================
================================================================================
================================================================================
Name                                    : Ten1Vrf2
Tenant                                  : Ten1
Routing Type                            : centralized
Centralized Routers                     : 10.64.196.56
                        : 10.64.196.57
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 MPLS L3 Handoff                  : true
Enable Graceful Restart                 : false
Enable NextHop Recursion                : false
Enable Default Information Originate    : false
EVPN Route Target                       : import 103:103
                                        : export 103:103
Non-EVPN Route Target                   : import 201:201
                                        : export 201:201
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

================================================================================
================================================================================

efa tenant vrf show --tenant Ten2 --detail
================================================================================
================================================================================
Name                                    : Ten2Vrf1
```

```
Tenant                                  : Ten2
Routing Type                            : centralized
Centralized Routers                     : 10.64.196.56
                          : 10.64.196.57
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 MPLS L3 Handoff                  : false
Enable Graceful Restart                 : false
Enable NextHop Recursion                : false
Enable Default Information Originate    : false
EVPN Route Target                       : import 102:102
                                        : export 102:102
Non-EVPN Route Target                   :
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

================================================================================
================================================================================
================================================================================
================================================================================
Name                                    : Ten2Vrf2
Tenant                                  : Ten2
Routing Type                            : centralized
Centralized Routers                     : 10.64.196.56
                          : 10.64.196.57
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 MPLS L3 Handoff                  : true
Enable Graceful Restart                 : false
Enable NextHop Recursion                : false
Enable Default Information Originate    : false
EVPN Route Target                       :
Non-EVPN Route Target                   : import 202:202
                                        : export 202:202
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

================================================================================
================================================================================



efa tenant vrf show --tenant Ten2
+----------+--------+--------------+---------------------+--------------
+--------------+----------+-----------+-----------------+-----------
+--------------------+-------------+-------------+
|   Name   | Tenant | Routing Type | Centralized Routers |  Enable L3   |
Redistribute | Max Path | Local Asn | MPLS L3 Handoff | Enable GR |
State        |  Dev State  |  App State  |
|          |        |              |                      |  Extension
|               |          |           |                 |
|                    |             |             |
+----------+--------+--------------+---------------------+--------------
+--------------+----------+-----------+-----------------+-----------
+--------------------+-------------+-------------+
| Ten2Vrf1 |  Ten2  | centralized  |    10.64.196.53     |    false     |
connected   |    8     |           |      false      |   false   | vrf-device-
created | provisioned | cfg-in-sync |
|          |        |              |    10.64.196.54     |
|               |          |           |                 |
|                    |             |             |
+----------+--------+--------------+---------------------+--------------
+--------------+----------+-----------+-----------------+-----------
+--------------------+-------------+-------------+
| Ten2Vrf2 |  Ten2  | centralized  |    10.64.196.53     |    false     |
connected   |    8     |           |      true       |   false   | vrf-device-
created | provisioned | cfg-in-sync |
|          |        |              |    10.64.196.54     |
|               |          |           |                 |
|                    |             |             |
+----------+--------+--------------+---------------------+--------------
+--------------+----------+-----------+-----------------+-----------
+--------------------+-------------+-------------+
Vrf Details
```

b. Verify the following configurations on SLX devices:

BL1

```
vrf Ten1Vrf1
 rd 172.31.254.190:1
 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
 !
!
vrf Ten2Vrf1
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
```

BL2

```
vrf Ten1Vrf1
 rd 172.31.254.190:1
 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
 !
!
vrf Ten2Vrf1
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
```

```
 !
 address-family ipv6 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
 !
!
vrf Ten1Vrf2
 rd 172.31.254.190:1
 address-family ipv4 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
!
vrf Ten2Vrf2
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
!
```

```
 !
 address-family ipv6 unicast
  route-target export 102:102 evpn
  route-target import 102:102 evpn
 !
!
vrf Ten1Vrf2
 rd 172.31.254.190:1
 address-family ipv4 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 201:201
  route-target import 201:201
  bgp next-hop loopback 3
 !
!
vrf Ten2Vrf2
 rd 172.31.254.190:10
 address-family ipv4 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
 address-family ipv6 unicast
  route-target export 103:103 evpn
  route-target import 103:103 evpn
  route-target export 202:202
  route-target import 202:202
  bgp next-hop loopback 3
 !
!
```

c. Validation Error: The following validation errors occur:

Command	Error
efa tenant vrf create --tenant "Ten1" --name "Ten1Vrf3" --routing-type "centralized" --layer3-extension-enable false --centralized-router "10.64.196.56,10.64.196.57" --max-path 8 --redistribute connected --non-evpn-rt 200:200 --non-evpn-rt-type both	Non-EVPN RTs not allowed when MPLS L3 Handoff is disabled
efa tenant vrf update --tenant "Ten1" --name "Ten1Vrf1" --operation mpls-l3-handoff-update --mpls-l3-handoff true	MPLS L3 Handoff Update not applicable for VRFs on devices
efa tenant vrf create --name default-vrf --tenant Ten1 --mpls-l3-handoff true	MPLS L3 Handoff can't be enabled/disabled for default-vrf

Command	Error
efa tenant epg create --tenant "Ten1" --name Ten1Epg3 --switchport-mode trunk --port 10.64.196.56[0/3] --vrf "Ten1Vrf3" --ctag-range 10 --anycast-ip 10:20.10.10.10/24	MPLS not enabled on Fabric fab1 for VRF Ten1Vrf3
efa tenant vrf create --tenant "t1" --name "vrf1" --routing-type "centralized" --layer3-extension-enable false --centralized-router "10.64.196.56,10.64.196.57" --max-path 8 --redistribute connected --mpls-l3-handoff true	MPLS not enabled on fabric fab1
efa tenant vrf create --tenant "t1" --name "vrf1" --layer3-extension-enable true --max-path 8 --redistribute connected --non-evpn-rt 200:200 --non-evpn-rt-type both --mpls-l3-handoff true	MPLS not enabled on fabric fab1
efa tenant vrf create --tenant "t1" --name "vrf1" --layer3-extension-enable true --max-path 8 --redistribute connected --mpls-l3-handoff true --non-evpn-rt 200:200 --non-evpn-rt-type both	MPLS not supported on device 10.64.196.53 (BR-SLX9540)

9. Configure MPLS QoS profile.
   a. To configure MPLS QoS profiles and bind them to devices or tenants, run the following command:

```
efa policy qos map create --type exp-tc-map --name EXPToTCMap1
            --rule "exp[0],tc[4],dp[1]" --rule "exp[4],tc[3],dp[0]" --rule
"exp[5],tc[1],dp[1]"
efa policy qos map create --type dscp-exp-map --name DSCPToEXPMap1
            --rule "dscp[11],exp[5]" --rule "dscp[34],exp[0]"
efa policy qos map create --type exp-dscp-map --name EXPToDSCPMap1
             -rule "exp[5],dscp[20]" --rule "exp[6],dscp[41]"
efa policy qos profile create --name QoSProfile1 --trust auto
            --exp-tc EXPToTCMap1 --dscp-exp DSCPToEXPMap1 --exp-dscp EXPToDSCPMap1

efa policy qos profile bind --name QoSProfile1 --device 10.64.164.45,10.64.164.44
```

b. Verify the following configuration on SLX devices:

BL1	**BL2**
qos-mpls map exp-traffic-class EXPToTCMap1 exp 0 to traffic-class 4 drop-precedence 1 exp 4 to traffic-class 3 drop-precedence 0 exp 5 to traffic-class 1 drop-precedence 1 ! qos-mpls map dscp-exp DSCPToEXPMap1 dscp 11 to exp 5 dscp 34 to exp 0 ! qos-mpls map exp-dscp EXPToDSCPMap1 exp 5 to dscp 20 exp 6 to dscp 41 ! qos-mpls map-apply exp-traffic-class EXP2TC2 All qos-mpls map-apply dscp-exp dscp2expmap2 All qos-mpls map-apply exp-dscp EXP2DSCP2 All	qos-mpls map exp-traffic-class EXPToTCMap1 exp 0 to traffic-class 4 drop-precedence 1 exp 4 to traffic-class 3 drop-precedence 0 exp 5 to traffic-class 1 drop-precedence 1 ! qos-mpls map dscp-exp DSCPToEXPMap1 dscp 11 to exp 5 dscp 34 to exp 0 ! qos-mpls map exp-dscp EXPToDSCPMap1 exp 5 to dscp 20 exp 6 to dscp 41 ! qos-mpls map-apply exp-traffic-class EXP2TC2 All qos-mpls map-apply dscp-exp dscp2expmap2 All qos-mpls map-apply exp-dscp EXP2DSCP2 All

c. Configure the QoS MPLS maps binding types.

i. Fabric Binding: The QOS MPLS maps bind globally to all leaf or border devices in fabric (not per interface). If the device version lacks MPLS support, it's skipped, binding still shows success. Use the following command:

```
efa policy qos profile bind --name newprofile2  --fabric "ncls-fab-test"
Efa policy qos profile bind -name newprofile2 -fabric "ncls-fab-test" --port
fabric-internal
```

ii. Tenant Binding: QOS MPLS maps apply globally on 10.64.196.63 (tenant's device), not port-specific. If the device version lacks MPLS support, it's skipped, binding still shows success. Use the following command:

```
efa tenant create --name "tenant4" --type private --vlan-range 165-190 --vrf-
count 0 --enable-bd  --port 10.64.196.63[0/4-5]
efa policy qos profile bind --name mplsQosPro2  --tenant tenant4
efa policy qos profile bind --name mplsQosPro2 --port 10.64.196.63[0/4] --tenant
tenant4
efa policy qos profile bind --name mplsQosPro2  --tenant tenant4 -po po1
```

iii. Device Binding: The QOS MPLS maps will be bound to the device specified at the global level. The QOS maps will be bound to all the interface of the device specified. Use the following command:

```
efa policy qos profile bind
--name mplsQosProfile1 --device 10.64.196.63
```

Binding Priority: The binding priority determines which binding takes precedence when multiple bindings are applied, which is as follows:

i. Tenant interface
ii. Tenant
iii. Device

iv. Fabric interface
v. Fabric

10. Migrate VRFs when you upgrade to XCO 4.1.0 or higher.

    You can migrate your VRFs to L3VPN MPLS-based L3 Handoff. When upgrading to XCO 4.1.0 or higher, enable L3VPN MPLS-based L3 Handoff.

    a. Enable Router MPLS.

       Enable MPLS at the fabric level, configure MPLS fabric settings, and run **`efa fabric configure`** command to set up Router MPLS Loopback ID on necessary devices.

    b. Create L3VPN MPLS L3 Handoff.

       Create shared tenant, port-channel, default-vrf, L3 Handoff EPG, and BGP Peer for L3VPN MPLS-based L3 Handoff. Existing VRFs do not auto-migrate.

    c. Delete IPv4/v6 Unicast L3 Handoff.

       Delete previously configured IPv4/v6 unicast-based L3 Handoff EPGs and related configurations, as L3VPN MPLS-based L3 Handoff is sufficient for private tenant VRFs.

    d. Delete and re-create L3 Extension EPGs.

       Delete and re-create previously configured L3 Extension EPGs.

Administered Partial Success

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 139.

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 594.
- 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 125.
- 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 772.
- 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, 4.1.0*.

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>
```

APS Behavior of Tenant Configuration

Use this topic to know about the APS behavior of tenant configuration.

Existing behavior

Configuration or Deconfiguration is never attempted on admin down switching devices.

Target Devices

Devices on which the configuration is intended to be pushed.

Complete Failure Topology

- Topology having at least one single-homed device in admin down state or atleast one dual-homed device pair with both the devices in admin down state is a “complete failure topology”.
- Any Tenant CLI or REST API of create or delete nature attempted on the target devices having “complete failure topology” is rejected with an appropriate error and result in a complete failure. XCO does not have any configuration recipe prepared for this REST API or CLI as the entire request is rejected. For example, an EPG (endpoint group) create attempted on a single-homed device which is admin down state.

Complete Success Topology

- Topology with all the single-homed devices and all the dual-homed device pair in admin up state is a “complete success topology”.
- Any Tenant CLI or REST API of create or delete nature attempted on the target devices having “complete success topology” results in the configuration recipe preparation for all the target devices and the configuration is attempted on all the target devices as all the target devices are in “admin up” state.

Partial Success Topology

- Topology which is not a “complete failure topology” and have at least one dual-homed device pair with one of the device in admin down state and the other device in admin up state is a “partial success topology”.
- Any Tenant CLI or REST API of create nature attempted on the target devices having “partial success topology” will result in the configuration being attempted on the “admin up” devices and configuration not being attempted on the “admin down” devices. Even though configuration is not attempted for “admin down” devices, the configuration is treated as success for the “admin down” devices.

  Configuration recipe is prepared and persisted in XCO for all the devices, and the configuration is auto reconciled with the devices when the “admin down” devices transition to “admin up”.

  When the devices are in the “admin up” state, the XCO intended configuration synchronizes with the device configuration.

- Any Tenant CLI or REST API of delete nature attempted on the target devices having “partial success topology” results in deconfiguration attempted on the “admin up” devices, and deconfiguration not attempted on the “admin down” devices. The CLI or REST API operation fails with an appropriate error indicating that the deconfiguration not being attempted on the “admin down” devices.

  The reason being XCO does not want to leave stale configurations on the devices because if the stale configurations are left on the devices, then bringing the devices (having stale configurations) back into XCO is erroneous considering the full brownfield support is missing in XCO. You can retry the same CLI or REST API operation after the “admin down” devices transition to “admin up” state so that the deconfiguration is attempted on all the devices. You can always use “force” option available in REST API to forcefully delete the entities from XCO even in case of partial success topology.
- Drift between the XCO intended configuration and device config is shown in the **`efa tenant debug device drift`** CLI or REST API output and in the corresponding entity GET or SHOW output (for example, **`efa tenant epg show`** and **`efa tenant po show`** in the form of “app-state” and “dev-state”.
- XCO blocks the tenant reconciliation API and rest of the tenant APIs support partial success behavior.

APS: Pre-provisioning Support by Modifying the Target Device List to Include the MCT Neighbor

Scenario	Target Device List	Resultant Topology
Single Homed PO Create with PO member on Rack1Device2	Rack1Device1 Rack1Device2	Partial Success
Endpoint group (EPG) create with cluster edge port (CEP) member on Rack1Device2	Rack1Device1 Rack1Device2	Partial Success
BGP Peer Group create with the peer-group residing on Rack1Device2	Rack1Device1 Rack1Device2	Partial Success
BGP Peer Create with the static or dynamic BGP peers residing on Rack1Device2	Rack1Device1 Rack1Device2	Partial Success
VRF update with SR or SR-BFD residing on Rack1Device2	Rack1Device1 Rack1Device2	Partial Success

```
efa tenant show
+------+-------------+-------------+------------+-----------+-----------+---------
+----------------------+
| Name | L2VNI-Range | L3VNI-Range | VLAN-Range
| VRF-Count | Enable-BD |  Type   |        Ports         |
+------+-------------+-------------+------------+-----------+-----------+---------
+----------------------+
| ten1 |             |             | 11-20      |
10        | False     | private | 10.20.246.15[0/1-10] |
|      |             |             |            |
|           |         | 10.20.246.16[0/1-10] |
```

```
+------+-------------+-------------+------------+-----------+-----------+---------
+----------------------+
Tenant Details

efa inventory admin-state down --ip 10.20.246.15
AdminStateDown [success]
Admin State Down execution UUID: 6eaa1ebe-40fe-4628-8d5c-df11ffc4521e
execute the CLI to get details : efa inventory
admin-state detail --uuid 6eaa1ebe-40fe-4628-8d5c-df11ffc4521e

efa inventory admin-state detail --uuid 6eaa1ebe-40fe-4628-8d5c-df11ffc4521e
+---------------------------------+----------------------------------------+
|              NAME               |                VALUE                   |
+---------------------------------+----------------------------------------+
| UUID                            | 6eaa1ebe-40fe-4628-8d5c-df11ffc4521e |
+---------------------------------+----------------------------------------+
| Device IP                       | 10.20.246.15                           |
+---------------------------------+----------------------------------------+
| Admin State Action              | down                                   |
+---------------------------------+----------------------------------------+
| Status                          | success                                |
+---------------------------------+----------------------------------------+
| Fabric Status                   | success                                |
+---------------------------------+----------------------------------------+
| Tenant Status                   | success                                |
+---------------------------------+----------------------------------------+
| Maintenance Mode Enable Status | success                                |
+---------------------------------+----------------------------------------+
| Start Time                      | 2021-02-06 21:18:53 -0800 PST          |
+---------------------------------+----------------------------------------+
| Last Modified                   | 2021-02-06 21:19:59 -0800 PST          |
+---------------------------------+----------------------------------------+
| Duration                        | 1m5.517263907s                         |
+---------------------------------+----------------------------------------+
```

Behavior in XCO

```
efa tenant po create --name ten1po1 --tenant ten1 --port 10.20.246.15[0/1-2] --speed 10Gbps --
negotiation active

efa tenant epg create --name ten1epg1 --tenant ten1 --po ten1po1 --switchport-mode trunk --ctag-range
11-12 --anycast-ip 11:10.0.11.1/24 --anycast-ip 12:10.0.12.1/24 --vrf ten1vrf1

efa tenant service bgp peer create --name ten1bgppeer1 --tenant ten1 --ipv4-uc-nbr
10.20.246.15,ten1vrf1:10.0.0.0,65001

efa tenant service bgp peer-group create --name ten1bgppeergroup1 --tenant ten1 --pg-name
10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:65010

efa tenant po show --name ten1po1 --tenant ten1
+---------+--------+----+--------+-------------+----------+---------+---------------------+------------
+-----------------+-----------+
|  Name   | Tenant | ID | Speed  | Negotiation | Min Link |  Lacp   |        Ports        |   State
|    Dev State    | App State |
|         |        |    |        |             |  Count   | Timeout |                     |
|                 |           |
+---------+--------+----+--------+-------------+----------+---------+---------------------+------------
+-----------------+-----------+
| ten1po1 |  ten1  | 1  | 10Gbps |   active    |    1     |  long   | 10.20.246.15[0/1-2] | po-created
| not-provisioned | cfg-ready |
+---------+--------+----+--------+-------------+----------+---------+---------------------+------------
+-----------------+-----------+

efa tenant vrf show --name ten1vrf1 --tenant ten1
```

```
(efa:root)root@node-2:~# efa tenant vrf show --name ten1vrf1 --tenant ten1
+----------+--------+-------------+-----------+------------------+-----+------+-------+------------
+-----------------+-----------+
|   Name   | Tenant | Routing    | Centralized |   Redistribute   | Max | Local| Enable|   State
|    Dev State    | App State |
|          |        | Type       | Routers     |                  | Path| Asn  | GR    |
|                 |           |
+----------+--------+------------+-------------+------------------+-----+------+-------+------------
+-----------------+-----------+
| ten1vrf1 |  ten1  | distributed|             | connected,static | 50  | 65002| false | vrf-create |
not-provisioned | cfg-ready |
+----------+--------+------------+-------------+------------------+-----+------+-------+------------
+-----------------+-----------+


efa tenant epg show --detail
=============================================================================================================
=======================================================
Name          : ten1epg1
Tenant        : ten1
Description   :
Type          : extension
Ports         :
POs           :
              : unstable            : ten1po1
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12
              : vrf                 : ten1vrf1 [unstable]
              : l3-vni              : 8192
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 |             |         |                               | Tenant L3
Extended VLAN |             | False                | False               | not-provisioned | cfg-
ready |
+------+--------+-------------+-------------+---------+-------------------------------
+-------------------------+-------------+----------------------+--------------------
+-----------------+-----------+
| 12   | 12     | 10.0.12.1/24 |             |         |                               | Tenant L3
Extended VLAN |             | False                | False               | not-provisioned | cfg-
ready |
+------+--------+-------------+-------------+---------+-------------------------------
+-------------------------+-------------+----------------------+--------------------
+-----------------+-----------+

efa tenant service bgp peer-group show
=============================================================================================================
=======================================================
Name        : ten1bgppeergroup1
Tenant      : ten1
State       : bgp-pg-state-created
Description :
+--------------+-----------------+--------+--------+----------+-----+-----------+---------------
+------------------+-----------------+------------+
|  Device IP   | PEER-GROUP-NAME | REMOTE | BFD    | BFD      | BFD | BFD       | Next-Hop-Self |
Update-Source-IP |    Dev-state    | App-state |
```

```
|              |                 | ASN    | Enabled| Interval | Rx  | Multiplier|
|                  |                |           |
+--------------+-----------------+--------+--------+----------+-----+-----------+---------------
+------------------+-----------------+-----------+
| 10.20.246.15 | pg1             | 65010  | false  | 0        | 0   | 0         | false
|                  | not-provisioned | cfg-ready |
+--------------+-----------------+--------+--------+----------+-----+-----------+---------------
+------------------+-----------------+-----------+

efa tenant service bgp peer show
================================================================================================================
================================================================================
Name        : ten1bgppeer1
Tenant      : ten1
State       : bs-state-created
Description :
Static Peer:
+--------------+----------+------+---------+----------+--------+--------+---------+-----+-----------
+----------+----------+-----------------+-----------+
|  Device IP   |   VRF    | AFI  |  SAFI   | REMOTE   | REMOTE | BFD    | BFD     | BFD | BFD       |
Next Hop | Update   |    Dev-state    | App-state |
|              |          |      |         | IP       | ASN    | Enabled| Interval| Rx  | Multiplier|
Self     | Source IP|                 |           |
+--------------+----------+------+---------+----------+--------+--------+---------+-----+-----------
+----------+----------+-----------------+-----------+
| 10.20.246.15 | ten1vrf1 | ipv4 | unicast | 10.0.0.0 | 65001  | false  | 0       | 0   | 0         |
false    |          | not-provisioned | cfg-ready |
+--------------+----------+------+---------+----------+--------+--------+---------+-----+-----------
+----------+----------+-----------------+-----------+
Dynamic Peer:
+-----------+-----+-----+------+--------------+------------+--------------+-----------+-----------+
| Device IP | VRF | AFI | SAFI | Listen Range | Peer Group | Listen Limit | Dev-state | App-state |
+-----------+-----+-----+------+--------------+------------+--------------+-----------+-----------+
================================================================================================================
================================================================================

efa inventory admin-state up --ip 10.20.246.15
AdminStateUp [success]
Admin State Up started execution UUID: 0ae0db00-c0de-4d55-8410-6d67bca4ad65


efa inventory admin-state detail --uuid 0ae0db00-c0de-4d55-8410-6d67bca4ad65
+------------------------------+----------------------------------------+
|            NAME              |                VALUE                   |
+------------------------------+----------------------------------------+
| UUID                         | 0ae0db00-c0de-4d55-8410-6d67bca4ad65 |
+------------------------------+----------------------------------------+
| Device IP                    | 10.20.246.15                           |
+------------------------------+----------------------------------------+
| Admin State Action           | up                                     |
+------------------------------+----------------------------------------+
| Status                       | success                                |
+------------------------------+----------------------------------------+
| Fabric Status                | success                                |
+------------------------------+----------------------------------------+
| Tenant Status                | success                                |
+------------------------------+----------------------------------------+
| Drift and Reconcile id       | a3c987ea-f709-4f7c-8ae4-bc5c9481cc55 |
+------------------------------+----------------------------------------+
| Drift and Reconcile Status   | DR Completed                           |
+------------------------------+----------------------------------------+
| Start Time                   | 2021-02-06 21:39:09 -0800 PST        |
+------------------------------+----------------------------------------+
| Last Modified                | 2021-02-06 21:47:09 -0800 PST        |
```

```
+-----------------------------+----------------------------------------+
| Duration                    | 8m0.126957723s                         |
+-----------------------------+----------------------------------------+
--- Time Elapsed: 47.334754ms ---
(efa:root)root@node-2:~#

efa tenant vrf show --name ten1vrf1 --tenant ten1
+----------+--------+--------------+-------------+------------------+-----+-------+--------
+--------------------+-------------+-------------+
|   Name   | Tenant | Routing Type | Centralized |   Redistribute   | Max | Local | Enable |
State        |  Dev State  |  App State  |
|          |        |              | Routers     |                  | Path| Asn   | GR
|                    |             |             |
+----------+--------+--------------+-------------+------------------+-----+-------+--------
+--------------------+-------------+-------------+
| ten1vrf1 |  ten1  | distributed  |             | connected,static | 50  | 65002 | false  | vrf-
device-created | provisioned | cfg-in-sync |
+----------+--------+--------------+-------------+------------------+-----+-------+--------
+--------------------+-------------+-------------+

efa tenant po show
+---------+--------+----+--------------------------+--------+-------------+--------------
+---------------------+-------------+------------+-------------+--------------+
|  Name   | Tenant | ID |       Description        | Speed  | Negotiation | MinLinkCount |
Ports        | LacpTimeout |   State    |  Dev-State  |   App-State  |
+---------+--------+----+--------------------------+--------+-------------+--------------
+---------------------+-------------+------------+-------------+--------------+
| ten1po1 | ten1   | 1  | EFA Port-channel ten1po1 | 10Gbps | active      | 1            |
10.20.246.15[0/1-2] | long        | po-created | provisioned |  cfg-in-sync |
+---------+--------+----+--------------------------+--------+-------------+--------------
+---------------------+-------------+------------+-------------+--------------+

efa tenant epg show
================================================================================================================================
============================================================
Name          : ten1epg1
Tenant        : ten1
Description   :
Type          : extension
Ports         :
POs           :
              : unstable            : ten1po1
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12
              : vrf                 : ten1vrf1
              : l3-vni              : 8192
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 |              |         |                               | Tenant L3
Extended VLAN |             | False                 | False               | provisioned |  cfg-in-sync
|
+------+--------+--------------+--------------+---------+-------------------------------
+-------------------------+-------------+-----------------------+---------------------+-------------
+--------------+
| 12   | 12     | 10.0.12.1/24 |              |         |                               | Tenant L3
```

```
Extended VLAN |             | False                | False               | provisioned |  cfg-in-sync
|
+------+--------+--------------+--------------+---------+--------------------------------
+-------------------------+-------------+----------------------+---------------------+-------------
+--------------+
For 'unstable' entities, run 'efa tenant po/vrf show' for details
================================================================================================================
============================================================


efa tenant service bgp peer-group show
================================================================================================================
============================================================
Name        : ten1bgppeergroup1
Tenant      : ten1
State       : bgp-pg-state-created
Description :
+--------------+-----------------+------------+-------------+--------------+---------+----------------
+---------------+------------------+--------------+-------------+
|  Device IP   | PEER-GROUP-NAME | REMOTE ASN | BFD Enabled | BFD Interval | BFD Rx | BFD Multiplier |
Next-Hop-Self | Update-Source-IP |  Dev-state  |  App-state  |
+--------------+-----------------+------------+-------------+--------------+---------+----------------
+---------------+------------------+--------------+-------------+
| 10.20.246.15 | pg1             | 65010      | false       | 0            | 0      | 0              |
false         |                  | provisioned | cfg-in-sync |
+--------------+-----------------+------------+-------------+--------------+---------+----------------
+---------------+------------------+--------------+-------------+
================================================================================================================
============================================================

efa tenant service bgp peer show
================================================================================================================
============================================================================================
Name        : ten1bgppeer1
Tenant      : ten1
State       : bs-state-created
Description :
Static Peer:
+--------------+----------+------+---------+-----------+------------+-------------+--------------
+--------+----------------+---------------+-----------------+-------------+-------------+
|  Device IP   |   VRF    | AFI  |  SAFI   | REMOTE IP | REMOTE ASN | BFD Enabled | BFD Interval | BFD
Rx | BFD Multiplier | Next Hop Self | Update Source IP |  Dev-state  |  App-state  |
+--------------+----------+------+---------+-----------+------------+-------------+--------------
+--------+----------------+---------------+-----------------+-------------+-------------+
| 10.20.246.15 | ten1vrf1 | ipv4 | unicast | 10.0.0.0  | 65001      | false       | 0            |
0      | 0              | false         |                 | provisioned | cfg-in-sync |
+--------------+----------+------+---------+-----------+------------+-------------+--------------
+--------+----------------+---------------+-----------------+-------------+-------------+
Dynamic Peer:
+-----------+-----+-----+------+--------------+------------+--------------+-----------+-----------+
| Device IP | VRF | AFI | SAFI | Listen Range | Peer Group | Listen Limit | Dev-state | App-state |
+-----------+-----+-----+------+--------------+------------+--------------+-----------+-----------+
================================================================================================================
============================================================================================
```

Note

During PST (Partial Success Topology), attempting to delete a tenant entity may move its dev-state to 'not-provisioned' and its app-state to 'cfg-ready/cfg-refreshed' state.

APS: Deletion Support for Pre-provisioned Configurations

Issue in XCO

Creation of XCO entities (PO and EPG) on an admin down device followed by the deletion of same EFA entities (PO and EPG) on the same admin down device used to fail even though the configuration was never pushed to the devices.

Solution

Succeed the deletion of the XCO entities (PO and EPG) if the resultant configuration to be deleted has never been pushed to the devices.

Pre-provisioned config

The pre-provisioned config is present in XCO DB and not present on the SLX.

```
efa inventory admin-state  show --
ip 10.20.246.15
+----------------------
+--------------+
|          NAME
|    VALUE     |
+----------------------
+--------------+
| Device IP
| 10.20.246.15 |
+----------------------
+--------------+
| Admin State
| down         |
+----------------------
+--------------+
| Health Check
Status  | Disable      |
+----------------------
+--------------+
```

```
efa inventory admin-state  show --
ip 10.20.246.16
+----------------------
+--------------+
|          NAME
|    VALUE     |
+----------------------
+--------------+
| Device IP
| 10.20.246.16 |
+----------------------
+--------------+
| Admin State
| up           |
+----------------------
+--------------+
| Health Check
Status  | Disable      |
+----------------------
+--------------+
```

```
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
efa tenant vrf create --name ten1vrf1 --tenant ten1
efa tenant epg create --name ten1epg1 --tenant ten1 --po ten1po1 --switchport-mode trunk
--ctag-range 11-12 --anycast-ip 11:10.0.11.1/24 --anycast-ip 12:10.0.12.1/24 --vrf
ten1vrf1
efa tenant service bgp peer create --name ten1bgppeer1 --tenant ten1 --ipv4-uc-nbr
10.20.246.15,ten1vrf1:10.0.0.0,65001 --ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.1.0.0,65001
efa tenant service bgp peer-group create --name ten1bgppeergroup1 --tenant ten1 --pg-name
10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:65010  --pg-name 10.20.246.16:pg1 --pg-asn
10.20.246.16,pg1:65010

efa tenant po show
+---------+--------+----+--------------------------+--------+-------------+--------------
+-------------------+-------------+------------+-----------------+-----------+
|  Name   | Tenant | ID |       Description        | Speed  | Negotiation | MinLinkCount
|       Ports       | LacpTimeout |   State    |    Dev-State    | App-State |
+---------+--------+----+--------------------------+--------+-------------+--------------
+-------------------+-------------+------------+-----------------+-----------+
| ten1po1 | ten1   | 1  | EFA Port-channel ten1po1 | 10Gbps | active      | 1
| 10.20.246.15[0/1] | long        | po-created | not-provisioned | cfg-ready |
|         |        |    |                          |        |             |
| 10.20.246.16[0/1] |             |            |                 |           |
```

```
+---------+--------+----+-------------------------+--------+-------------+--------------
+-------------------+-------------+-----------+-----------------+-----------+

efa tenant service bgp peer-group show
================================================================================================
=================================================================
Name        : ten1bgppeergroup1
Tenant      : ten1
State       : bgp-pg-state-created
Description :
+--------------+----------------+-----------+------------+--------------+--------
+---------------+--------------+-----------------+-----------------+-------------+
|  Device IP   | PEER-GROUP-NAME | REMOTE ASN | BFD Enabled | BFD Interval | BFD Rx | BFD
Multiplier | Next-Hop-Self | Update-Source-IP |    Dev-state    |  App-state  |
+--------------+----------------+-----------+------------+--------------+--------
+---------------+--------------+-----------------+-----------------+-------------+
| 10.20.246.16 | pg1             | 65010      | false       | 0            | 0      |
0              | false         |                  | provisioned     | cfg-in-sync |
+--------------+----------------+-----------+------------+--------------+--------
+---------------+--------------+-----------------+-----------------+-------------+
| 10.20.246.15 | pg1             | 65010      | false       | 0            | 0      |
0              | false         |                  | not-provisioned | cfg-ready   |
+--------------+----------------+-----------+------------+--------------+--------
+---------------+--------------+-----------------+-----------------+-------------+
================================================================================================
=================================================================

efa tenant vrf show --name ten1vrf1 --tenant ten1
================================================================================================
=================================================================
Name                    : ten1vrf1
Vrf State               : vrf-device-created
Vrf Device State        : not-provisioned
Vrf App State           : cfg-ready
Tenant Name             : ten1
Routing Type            : distributed
L3 VNI                  : 8191
IRB BD                  : 4095
IRB VE                  : 8191
BR BD                   :
BR VE                   :
BR VNI                  : 4096
RH max path             :
RH ecmp enable          :
Graceful restart enable :
Route Target            : import 101:101
                        : export 101:101
Static Route            : Switch-IP->Network,Nexthop-IP[Route-Distance], ...
                        :
Local Asn               :
Static Route BFD        : Switch-IP->[DestIP,SourceIP][Interval,Min-Rx,Multiplier], ...
                        :
Max Path                : 8
Redistribute            : connected
================================================================================================
=================================================================

efa tenant epg show
================================================================================================
=================================================================
Name          : ten1epg1
Tenant        : ten1
Description   :
Type          : extension
```

```
Ports         :
POs           :
              : unstable            : ten1po1
Port Property : switchport mode     : trunk
              : native-vlan-tagging : false
NW Policy     : ctag-range          : 11-12
              : vrf                 : ten1vrf1 [unstable]
              : l3-vni              : 8191
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 |              |         |
| Tenant L3 Extended VLAN |             | False                 | False               |
not-provisioned | cfg-ready |
+------+--------+--------------+--------------+---------+--------------------------------
+-------------------------+-------------+-----------------------+---------------------
+-----------------+-----------+
| 12   | 12     | 10.0.12.1/24 |              |         |
| Tenant L3 Extended VLAN |             | False                 | False               |
not-provisioned | cfg-ready |
+------+--------+--------------+--------------+---------+--------------------------------
+-------------------------+-------------+-----------------------+---------------------
+-----------------+-----------+
For 'unstable' entities, run 'efa tenant po/vrf show' for details
================================================================================================
=====================================================================

efa tenant service bgp peer show
================================================================================================
================================================================================================
========
Name        : ten1bgppeer1
Tenant      : ten1
State       : bs-state-created
Description :
Static Peer:
+--------------+----------+------+---------+-----------+------------+-------------
+--------------+--------+----------------+---------------+------------------
+-----------------+-------------+
|  Device IP   |   VRF    | AFI  |  SAFI   | REMOTE IP | REMOTE ASN | BFD Enabled | BFD
Interval | BFD Rx | BFD Multiplier | Next Hop Self | Update Source IP |    Dev-state
|  App-state  |
+--------------+----------+------+---------+-----------+------------+-------------
+--------------+--------+----------------+---------------+------------------
+-----------------+-------------+
| 10.20.246.16 | ten1vrf1 | ipv4 | unicast | 10.1.0.0  | 65001      | false       |
0            | 0      | 0              | false         |                  |
provisioned     | cfg-in-sync |
+--------------+----------+------+---------+-----------+------------+-------------
+--------------+--------+----------------+---------------+------------------
+-----------------+-------------+
| 10.20.246.15 | ten1vrf1 | ipv4 | unicast | 10.0.0.0  | 65001      | false       |
0            | 0      | 0              | false         |                  | not-
provisioned | cfg-ready   |
+--------------+----------+------+---------+-----------+------------+-------------
+--------------+--------+----------------+---------------+------------------
+-----------------+-------------+
```

```
Dynamic Peer:
+-----------+-----+-----+------+--------------+-----------+-------------+-----------
+-----------+
| Device IP | VRF | AFI | SAFI | Listen Range | Peer Group | Listen Limit | Dev-state |
App-state |
+-----------+-----+-----+------+--------------+-----------+-------------+-----------
+-----------+
================================================================================================
================================================================================================
========
```

Scenario (Deletion of pre-provisioned config)	XCO
`efa tenant service bgp peer delete --name ten1bgppeer1 --tenant ten1`	`BgpService deleted successfully`
`efa tenant service bgp peer-group delete --name ten1bgppeergroup1 --tenant ten1`	`BgpService deleted successfully`
`efa tenant epg delete --name ten1epg1 --tenant ten1`	`EndpointGroup: ten1epg1 deleted successfully`
`efa tenant po delete --name ten1po1 --tenant ten1`	`PortChannel: ten1po1 deleted successfully`

Traffic Mirroring

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.
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, 4.1.0*.

In-band Traffic Mirroring

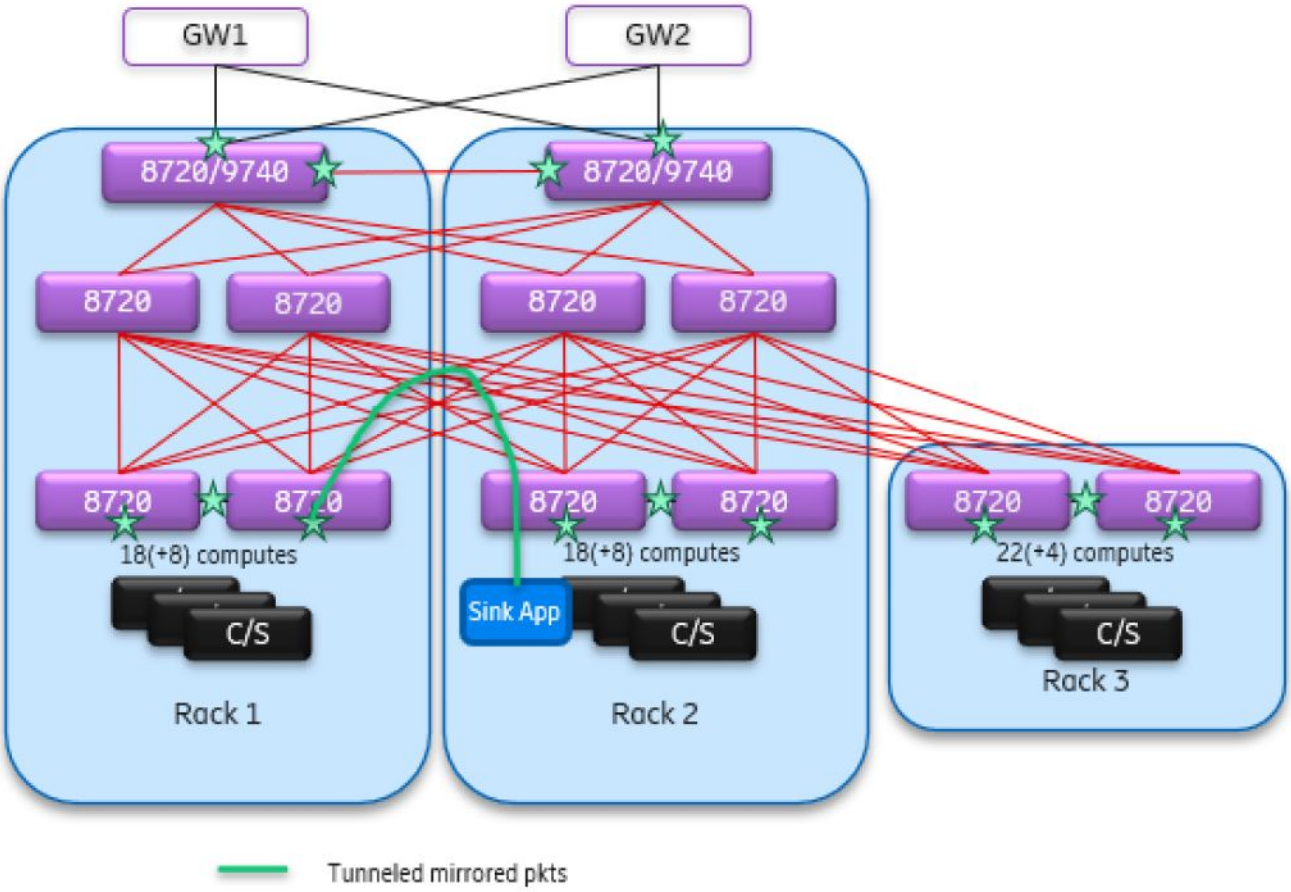

Figure 26: In-band traffic mirroring topology

- 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

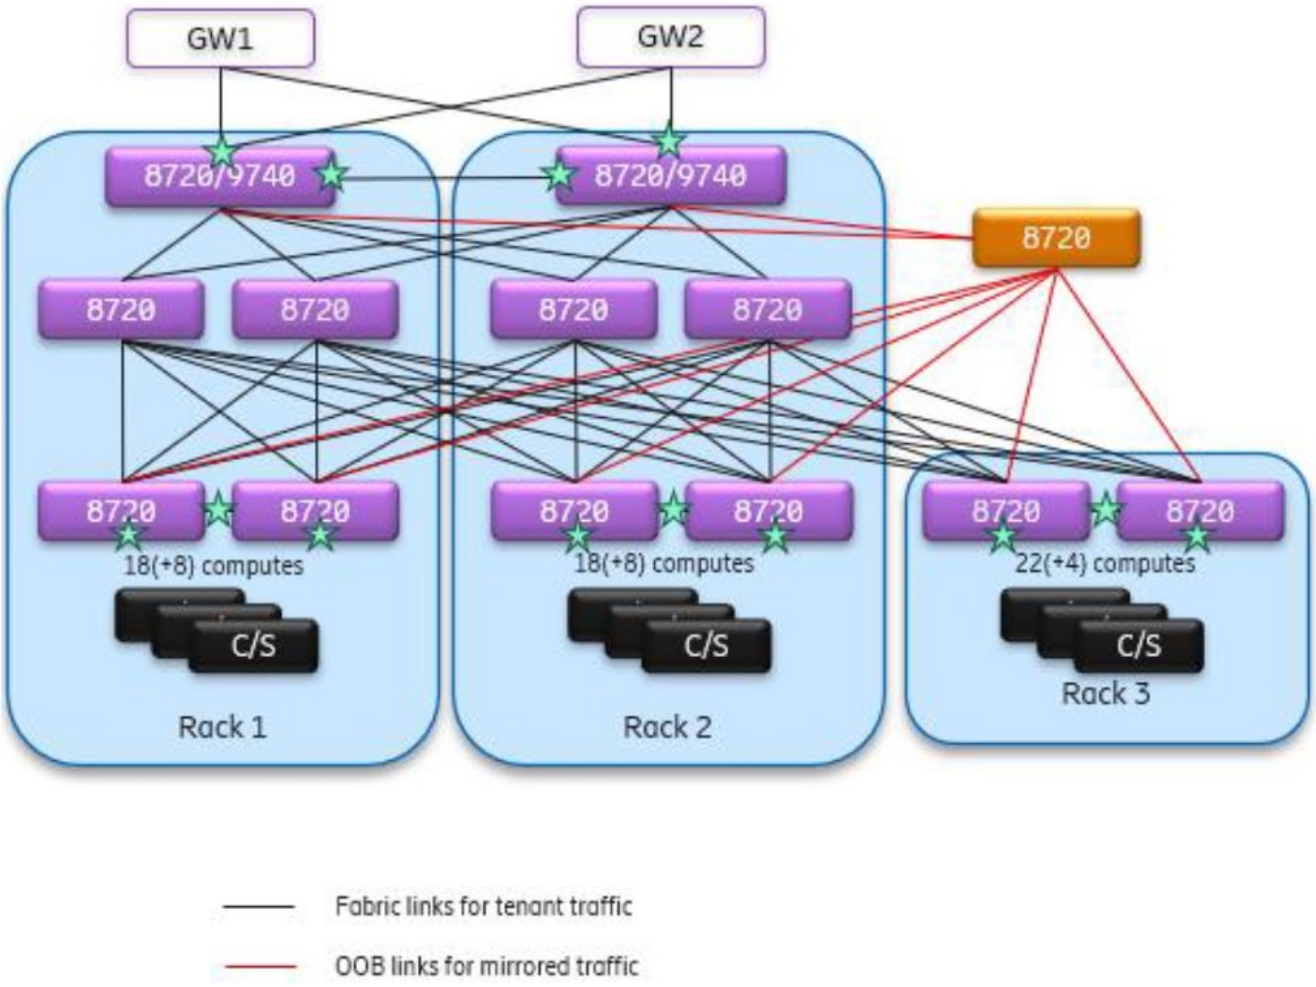

Figure 27: Out-of-band traffic mirroring topology

- 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

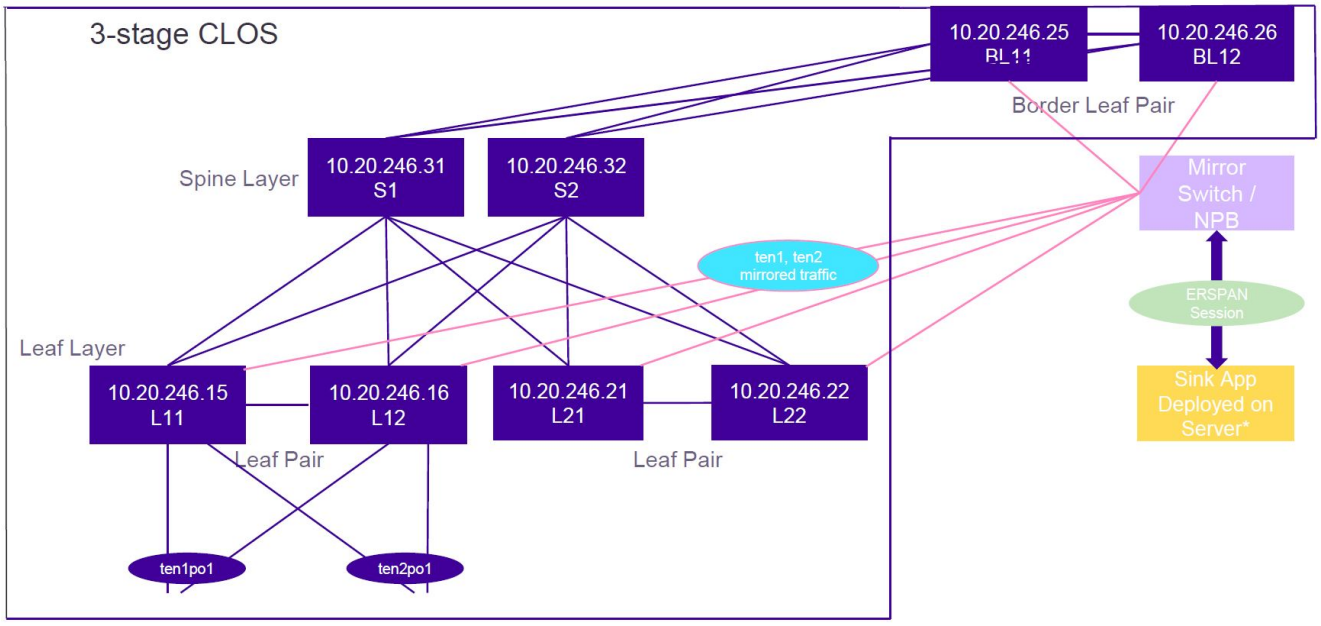

Figure 28: Port-based traffic mirroring topology

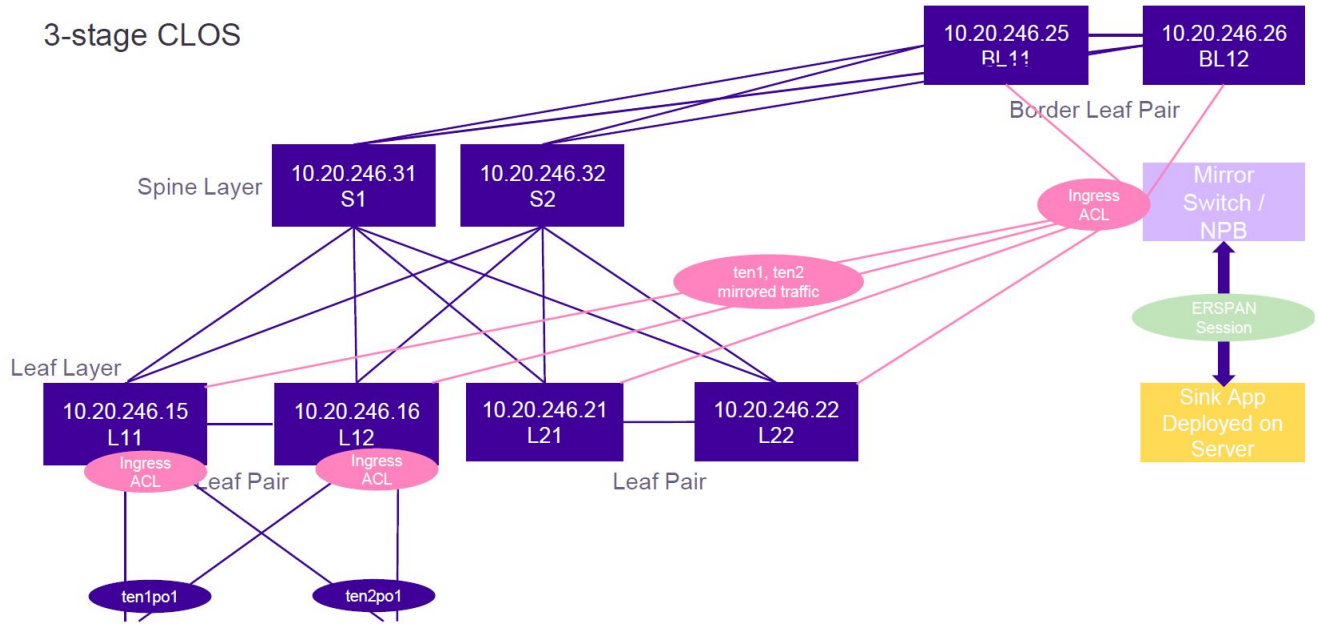

Figure 29: Flow-based traffic mirroring topology

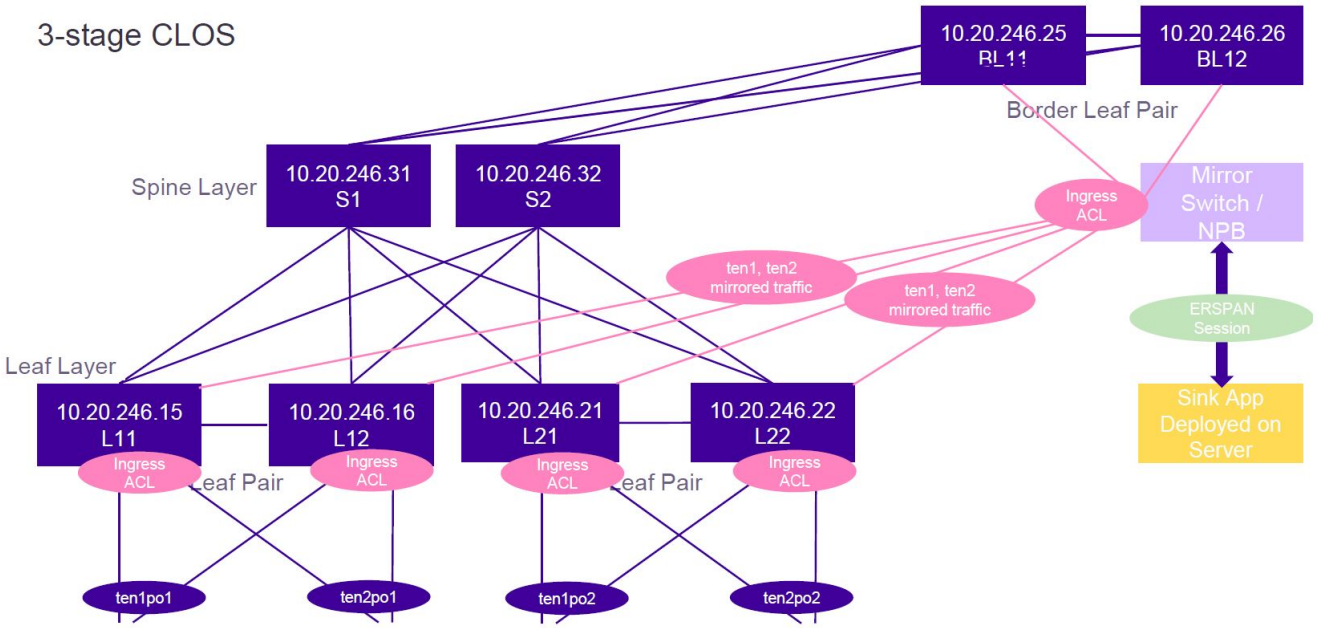

Figure 30: VLAN-based traffic mirroring topology

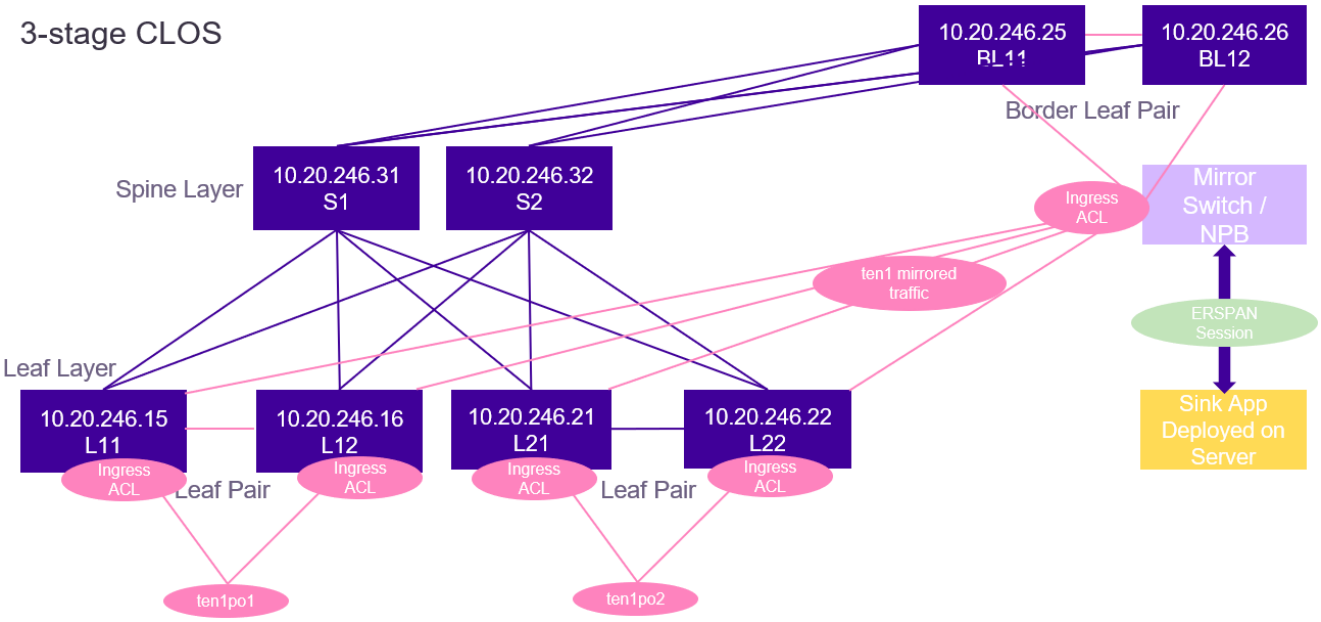

Figure 31: ICL port traffic mirroring topology

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 612
2. Configure Flow-Based Mirroring in a Multi-Tenant Architecture on page 614
3. Configure VLAN-Based Mirroring in a Multi-Tenant Architecture on page 618
4. Configure ICL Port Mirroring in a Multi-Tenant Architecture on page 622
5. Configure Fabric Non-ICL Ports as Mirror Source on page 626

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

```
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
```

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
```

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 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,ten1po1:10.20.246.1
5,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 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,ten1po1:10.20.246.1
5,eth,0/31
  --destination-type
10.20.246.15,po,ten1po1:span
  --direction
10.20.246.15,po,ten1po1:rx
```

```
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 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,ten1po2:10.20.246.2
1,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 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,ten1po2:10.20.246.2
1,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.

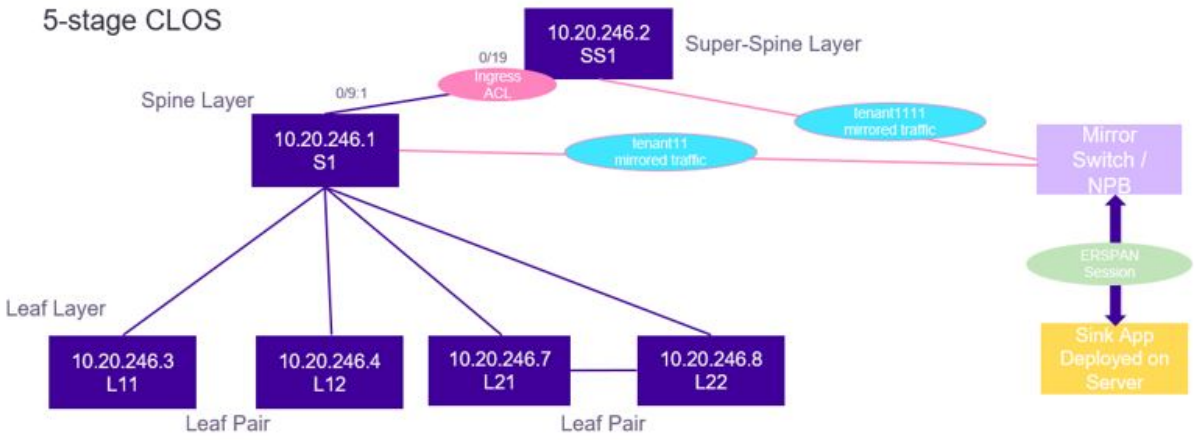

Figure 32: 5-stage Clos topology

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
```

Delete Pending Mirror Session Configuration

You can delete pending mirror session configuration.

About This Task

Follow this procedure to remove the pending mirror session configuration.

Procedure

Run the following command:

```
efa tenant service mirror session configure
```

The **`efa tenant service mirror session configure`** command pushes or removes a pending configuration for a mirror session when it is in `mirror-session-delete-pending` state.

Example

```
efa tenant service mirror session show
================================================================================================================
=====================================================================
Name        : m1
Tenant      : tv3
State       : mirror-session-delete-pending
Description :

+------+------------+---------------------------+-------------
+---------------------------+-----------------------+-----------+-------------
+-------------+
| Name |    Type    |          Source           | Destination |
Destination          |   Device Session ID    | Direction |  Dev State  |  App State  |
|      |            | [Device-IP,IfType,IfName]  |     Type    | [Device-
IP,IfType,IfName]  | [Device-IP,SessionID]  |           |             |             |
+------+------------+---------------------------+-------------
+---------------------------+-----------------------+-----------+-------------
+-------------+
|  m1  | port-based |    10.20.61.90,eth,0/6    |    span     |
10.20.61.90,eth,0/11    |     10.20.61.90,1      |   both    | provisioned | cfg-in-sync |
|      |            |                           |
|                           |                       |           |
|             |
+------+------------+---------------------------+-------------
+---------------------------+-----------------------+-----------+-------------
+-------------+

================================================================================================================
=====================================================================

--- Time Elapsed: 371.329541ms ---


efa tenant service mirror session configure --name m1 --tenant tv3

Mirror Service Session configured successfully.


--- Time Elapsed: 6.232577569s ---
(efa:extreme)extreme@node-1:~$ efa tenant service mirror session show

--- Time Elapsed: 286.058769ms ---
```

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

Learn about configuring policy service, such as IP prefix list, community list, route maps, and QoS support on fabric devices.

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 or IPv6 prefix lists on a set of fabric devices. If you have not specified the ge or le value, then the 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, 4.1.0*.

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                    |
+-----------+--------+-----+-------------------+--------------------------------+
```

Prefix List for Extreme OS ONE Devices

Learn about the prefix-list support for Extreme OS ONE devices in Extreme Cloud Orchestrator (XCO).

The enhancement extends existing XCO prefix-list functionality to Extreme OS ONE platforms and provides a unified CLI workflow for creating, updating, deleting, and listing prefix-lists across supported device families.

A prefix-list is a named collection of IPv4 or IPv6 prefixes used for route filtering operations. Extreme OS ONE implements prefix filtering through route-policy prefix-set constructs, while XCO provides a normalized CLI experience for operational consistency.

The prefix-list support to Extreme OS ONE platforms facilitate:

- Centralized management of network filtering rules across multiple devices
- Support for both IPv4 and IPv6 networks
- Flexible prefix matching with mask ranges (exact match or range-based)
- Simplified route filtering configuration

Supported Platforms: This feature is available for Extreme OS ONE devices managed through XCO.

Supported Capabilities

- Create prefix-lists for Extreme OS ONE devices
- Update existing prefix-lists
- Delete prefix-lists
- List configured prefix-lists
- Configure IPv4 and IPv6 prefixes
- Support ge and le mask-range qualifiers
- Synchronize prefix-list configurations with managed devices

Unsupported Capabilities

The following SLX-supported attributes are not applicable to Extreme OS ONE devices:

- Sequence number (seq)
- Action (permit or deny)

In Extreme OS ONE:

- Prefix-sets are match-only objects.
- Permit or deny actions are defined in route-policies.
- Idempotency is not supported in the policy service.
- AST topology is not supported.
- In-flight transactions are not supported.
- Native sequence number, type, and action are not supported on OS ONE. XCO accepts these parameters for compatibility but ignores them when configuring OS ONE devices.

A unified prefix-set supports both IPv4 and IPv6 entries. This feature is available for Extreme OS ONE devices managed through XCO.

Prefix List Behavioral Differences Between SLX and Extreme OS ONE

To support prefix lists on Extreme OS ONE, which use a different parameter model and behavior from SLX, XCO introduces the following changes to maintain a consistent northbound interface across both platforms:

- **Action**, **Sequence Number**, and **Type** are now optional when creating a prefix list.
- Prefix lists created without these optional parameters can be applied only to Extreme OS ONE devices. Attempting to apply such prefix lists to SLX devices results in a validation error.
- For backward compatibility, when a prefix list containing all optional parameters is bound to an Extreme OS ONE device, XCO ignores those parameters during device configuration.
- After the prefix list is bound, XCO automatically creates an additional **extrapolated** prefix list entry in the command output. This entry is a copy of the user-defined prefix list with the optional parameters removed, matching the format configured on the Extreme OS ONE device. The **extrapolated** flag is set to **true** for this entry.
- Extrapolated prefix list entries are system-generated and cannot be modified or deleted. These entries also cannot be referenced in bind operations for other devices.

Note
Migration of prefix lists from SLX to Extreme OS ONE is not supported.

Functionality	SLX behavior	Extreme OS ONE behavior
Sequence Number	Each entry has a sequence number.	No sequence number is configured.
Action (permit/ deny)	The action (permit or deny) is part of the prefix list.	The prefix set is match-only; action, sequence number, and type come from the route policy (policy result).
Type (IPv4/IPv6)	Separate commands exist for IPv4 and IPv6 (ip prefix-list / ipv6 prefix-list).	A single prefix set supports both IPv4 and IPv6.
Prefix range (ge / le)	The ge and le qualifiers are optional.	A mask-range is mandatory, specified as an exact value or as a range (X..Y).

Prefix Range Conversion

On SLX, you express prefix ranges with the ge and le qualifiers. On Extreme OS ONE, the equivalent range is expressed as a mask-range. XCO converts SLX-style ranges to Extreme OS ONE mask-ranges as shown in the following table:

SLX syntax	Extreme OS ONE syntax
10.0.0.0/8	10.0.0.0/8 mask-range exact
10.0.0.0/8 ge 16 le 24	10.0.0.0/8 mask-range 16..24

SLX syntax	Extreme OS ONE syntax
10.0.0.0/8 ge 20	10.0.0.0/8 mask-range 20..32 (IPv4)
10.0.0.0/8 le 24	10.0.0.0/8 mask-range 8..24
2001:db8::/32 ge 48 le 64	2001:db8::/32 mask-range 48..64
2001:db8::/32 ge 64	2001:db8::/32 mask-range 64..128 (IPv6)

Configuration Drift Behavior on Extreme OS ONE

On Extreme OS ONE, the `cfg-refreshed` state does not apply, because all rules are maintained under the same prefix list entry. If a prefix list contains one existing rule and you delete that rule, XCO marks the original rule as cfg-deleted and represents the newly added rule as `cfg-added`.

The policy service reports configuration drift as shown below:

Name	Type	Seq	App State	Child config
prefix_tos	prefix-list	0	cfg-entry-added	matchPrefixValue: prefix[10.0.0.0/8],ge[17],le[31]
prefix_tos	prefix-list	0	cfg-entry-deleted	matchPrefixValue: prefix[10.0.0.0/8],ge[16],le[30]

The corresponding device configuration change:

```
vtos121(config-route-policy-prefix-set-prefix_tos)# no prefix 10.0.0.0/8 mask-range 16..30
vtos121(config-route-policy-prefix-set-prefix_tos)# prefix 10.0.0.0/8 mask-range 17..31

route-policy
  prefix-set prefix_tos
    prefix 10.0.0.0/8 mask-range 17..31
```

Extreme OS ONE Route-Policy Representation

```
route-policy
prefix-set prefix1
prefix 10.1.1.0/24 mask-range exact
prefix 20.1.1.0/24 mask-range 25..30
```

Running Configuration Output

```
route-policy
prefix-set prefix1
   prefix 10.1.1.0/24 mask-range exact
   prefix 20.1.1.0/24 mask-range 25..30
```

Configure IP Prefix List on Extreme OS ONE Devices

You can configure IP prefix list on Extreme OS ONE devices.

About This Task

Follow this procedure to configure IP prefix list on Extreme OS ONE devices

Procedure

1. Create a prefix list.

```
efa policy prefix-list create
--name prefix_1_in
--rule "prefix[2001:db8::/48],ge[48],le[128]"

Name: prefix_1_in
+------+---------+--------+---------------+----+----+---------+
| Type | Seq num | Action | Prefix        | Ge | Le | Status  |
+------+---------+--------+---------------+----+----+---------+
|      |         |        | 2001:db8::/48 |48  |128 | Success |
+------+---------+--------+---------------+----+----+---------+
|      |         |        | 10.0.0.0/8    |18  | 30 | Success |
+------+---------+--------+---------------+----+----+---------+
```

Note

- IPv4 and IPv6 prefixes can coexist in the same prefix-list.
- If seq or action values are provided, they are ignored for Extreme OS ONE devices.
- Prefix validation is performed before configuration deployment.

2. Update a prefix list.

   Update or add more prefixes to an existing list.

```
efa policy prefix-list update --name customer_routes
    --operation add-rule
    --rule "prefix[10.0.0.0/8],ge[10],le[24]"
```

Note

- A prefix list created *`with`* type, sequence, and action requires the `--type` flag during update. A prefix list created *`without`* type, sequence, and action does not require the `--type` flag during update; if you provide it, XCO returns an error.
- During device binding, XCO internally ignores unsupported attributes and creates extrapolated entries on the Extreme OS ONE device.
- Extrapolated entries created during binding are expected and are not duplicate or unintended entries.
- If you attempt to add Extreme OS ONE only prefix list (one created without sequence, action, and type) to an SLX device, XCO returns an error: the selected prefix does not have sequence, action, and type, and is supported only for Extreme OS ONE devices.

The following is a sample prefix list after binding to both SLX and Extreme OS ONE devices:

Type	Seq	Action	Prefix	Ge	Le	Device IP	App State	Extrapolated
ipv6	11	permit	2001:db8::/48	48	128	10.20.246.29	cfg-in-sync	False
			2001:db8::/48	48	128	10.20.246.30	cfg-in-sync	True

In the output above, the entry with `Extrapolated = True` is the Extreme OS ONE-compatible copy that XCO generated automatically for the Extreme OS ONE device.

3. Apply prefix list to devices.
   The following command pushes the `customer_routes` prefix list to devices with IP addresses 10.20.246.29 and 10.20.246.30.
   ```
   efa policy prefix-list update --name customer_routes --operation add-device --ip
   10.20.246.29,10.20.246.30
   ```
4. View a prefix list.
   You can view all prefix lists on specific devices

   The following example shows a range (10.20.246.29-30), which includes both 10.20.246.29 and 10.20.246.30.
   ```
   efa policy prefix-list list --ip 10.20.246.29-30
   ```
5. Remove rules from a prefix list.
   Removes a specific prefix from the list. You must provide the exact prefix and mask range to remove.
   ```
   efa policy prefix-list update --name [list name] --operation remove-rule --rule
   "prefix[address],ge[min],le[max]"
   ```
6. Delete a prefix list.
   The command deletes the entire prefix list.
   ```
   efa policy prefix-list delete --name plist2

   +--------------+---------+--------+
   | IP Address   | Result  | Reason |
   +--------------+---------+--------+
   | 10.20.246.29 | Success |        |
   | 10.20.246.30 | Success |        |
   +--------------+---------+--------+
   ```

   Note
   - XCO validates the prefix-list name before deletion.
   - Prefix-lists associated with BGP peers or peer-groups cannot be deleted until bindings are removed.

7. List Prefix-Lists.
   ```
   efa policy prefix-list list --ip 10.20.246.29

   Name: prefix_1_in_v4_1 (Extrapolated from: prefix_1_in)
   Type*    Seq*    Action*    Prefix    Ge    Le    Device IP    App State
   Extrapolated
              10.0.0.0/8    18    32    10.20.246.29    cfg-in-sync    true

   +-------+------+---------+-------------+----+----+--------------+-------------
   +--------------+
   | Type* | Seq* | Action* | Prefix      | Ge |
   Le | Device IP    | App State   | Extrapolated |
   +-------+------+---------+-------------+----+----+--------------+-------------
   +--------------+
   |       |      |         | 10.0.0.0/8  |    |    |
   10.20.246.29 | cfg-in-sync | true         |
   +------+-------+---------+-------------+----+----+--------------+-------------
   +--------------+
   *  Not applicable for OS ONE devices.
   ```

** Not applicable for Extreme OS ONE devices.*

Note

- Extreme OS ONE does not maintain separate IPv4 or IPv6 prefix-list types.
- If `--type` is used for Extreme OS ONE

Example (without device filter):

```
efa policy prefix-list list
Prefix-list details:
Name: prefix_1_in
+------+---------+--------+---------------+----+----+
| Type | Seq num | Action | Prefix        | Ge | Le |
+------+---------+--------+---------------+----+----+
|      |         |        | 2001:db8::/48 |    |    |
+------+---------+--------+---------------+----+----+
```

Note

- Extreme OS ONE does not maintain separate IPv4 or IPv6 prefix-list types.
- If `--type` is used for Extreme OS ONE devices, matching prefixes are not displayed.

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.

Note

For information about commands and supported parameters to configure route map, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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
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	

Configure Route-Map Next-Hop

You can configure the next-hop when you create a route map.

About This Task

Starting from XCO 4.0.2, route-map allows the configuration of the 'next-hop <value>' option (--set-next-hop), which can be used to set the next hop for eBGP IPv6 prefixes advertised to a specific neighbor.

The next-hop IPv6 address configuration is supported only on the Extreme SLX platform.

Follow this procedure to configure the next-hop when you create a route map.

Procedure

1. To configure the next-hop, run the **route-map set** command:

```
# efa policy route-map-set create --set-next-hop ipv4|ipv6,next-hop-address
```

The following example configures the next-hop.

```
# efa policy route-map-set create --name foo --rule seq[10],action[permit] --set-
community 6550:125,internet,local-as --set-next-hop ipv6,2004:22
+----------------+---------+--------+
| 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 configuration on SLX devices.

```
SLX# show running-config route-map
route-map foo permit 10
set community 6550:125 local-as internet
set ipv6 next-hop 2004:22
```

2. To remove the next-hop, run the **route-map set** command:

```
# efa policy route-map-set delete --next-hop ipv6,next-hop-address
```

The following example removes the next-hop.

```
efa policy route-map-set delete --name foo --rule seq[10],action[permit] --set-
community 6550:125,internet,local-as --set-next-hop ipv6,2004:22
+----------------+---------+--------+
| 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 configuration on SLX devices.

```
SLX# show running-config route-map
route-map foo permit 10
!
```

3. To view information about the route-map configured and its application on devices, run the **route-map list** command.

```
efa policy route-map list --ip 10.20.246.29 --detail
Route Map details:
```

```
Name: rmap-oob
Seq: 32
Action: permit
Matches:
match: matchExtcommunityList eclist-oob [cfg-not-managed]
match: matchCommunityList clist-oob [cfg-not-managed]
match: matchLargeCommunityList lclist-oob [cfg-not-managed]
Sets:
set: setLargeCommunityList lclist-oob [cfg-not-managed]
Name: rmap1
Seq: 5
Action: permit
Matches:
match: matchExtcommunityList extcommunitystd1
Sets:
Name: rmap1
Seq: 65535
Action: permit
Matches:
Sets:
 set: setCommunityValue 6550:125 internet local-as setNextHopIpv6 2004::22
IP Addresses:
+----------+-------+--------------+-----------------+
| Name     | Seq   | IP Address   | App State       |
+----------+-------+--------------+-----------------+
| rmap-oob | 32    | 10.20.246.29 | cfg-not-managed |
+----------+-------+--------------+-----------------+
| rmap1    | 5     | 10.20.246.29 | cfg-in-sync     |
+----------+-------+--------------+-----------------+
| rmap1    | 65535 | 10.20.246.29 | cfg-in-sync     |
+----------+-------+--------------+-----------------+
```

Route Map for Extreme OS ONE Devices

Learn about route map integration capabilities supported by XCO for Extreme OS ONE devices. The implementation introduces support for route-map, route-map-match, and route-map-set operations through XCO-managed CLI workflows.

A route map provides a policy-based framework for controlling routing decisions. Each route-map statement includes:

- Match conditions that identify routes based on route attributes
- Set actions that modify selected route attributes
- Permit or deny logic that defines policy behavior for matching routes

Route maps support routing policy enforcement, traffic engineering, and route filtering across managed devices. This feature does not store personal user data in persistent storage. All operational data is cleared after a device reboot.

Supported Platforms: This feature is available for Extreme OS ONE devices managed through XCO.

Supported Capabilities

The following operations are supported:

- Route-map create, update, delete, and list operations
- Route-map synchronization with managed devices (Multi-device synchronization)
- Route-map-match configuration for:
  - IPv4 prefix lists
  - IPv6 prefix lists
  - Permit and deny actions
- Route-map-set configuration for:
  - Community list delete
  - Next-hop modification (IPv4 and IPv6)

Unsupported Capabilities

The following capabilities are not supported on E1OS devices:

Route-map-match

- Large community matching

Route-map-set

- Large community
- Large community delete
- Extended community SoO (Site of Origin)
- Extended community RT (Route Target) via `--set-extcommunity-rt` parameter
- Community inline values

Behavioral Differences Between SLX and Extreme OS ONE

Prefix Matching

Extreme OS ONE does not maintain separate configuration constructs for IPv4 and IPv6 prefix matching. Both address families use the same `match-prefix-set` configuration.

- SLX Example:

```
match ip address telco_1_in
match ipv6 address prefix-list telco_1_in
```

- Extreme OS ONE Example:

```
match-prefix-set telco_1_in any
```

Next-Hop Configuration

Extreme OS ONE uses a unified next-hop configuration model for both IPv4 and IPv6 addresses. For example,

```
set-next-hop 10.23.233.44
set-next-hop ff::ff
```

Community Handling

Large community operations are not supported on Extreme OS ONE devices. XCO displays guidance indicating that the feature is unsupported when applicable.

Extreme OS ONE-Specific Match Prefix Configuration

For Extreme OS ONE-only deployments that do not include SLX devices, XCO supports the `--match-prefix-set` token.

```
efa policy route-map-match create
--name rmap_1
--rule seq[10],action[permit]
--match-prefix-set prefix_v6
```

Note
The `--match-prefix-set` token is supported only for oneOS environments. SLX devices require IPv4 or IPv6 prefix type specification explicitly.

Device Configuration Examples

Create a Route Policy

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 10
Switch(config-route-policy-policy-map1-stmt-10)# exit
Switch(config-route-policy-policy-map1)# end
show running-config output
route-policy
      policy map1
          statement 10
```

Match Prefix Set

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 10
Switch(config-route-policy-policy-map1-stmt-10)# conditions
Switch(config-route-policy-policy-map1-stmt-10-conditions)# match-prefix-set p1 any
show running-config output
route-policy
        policy map1
           statement 10
               conditions
                   match-prefix-set p1 any
```

Match Community Set

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 20
Switch(config-route-policy-policy-map1-stmt-20)# conditions
Switch(config-route-policy-policy-map1-stmt-20-conditions)# bgp-conditions
Switch(config-route-policy-policy-map1-stmt-20-conditions-bgp-conditions)# match-
community-set comm1
show running-config output
route-policy
  policy map1
    statement 20
      conditions
        bgp-conditions
          match-community-set comm1
```

Match Ext-Community-Set

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 30
Switch(config-route-policy-policy-map1-stmt-30)# conditions
Switch(config-route-policy-policy-map1-stmt-30-conditions)# bgp-conditions
Switch(config-route-policy-policy-map1-stmt-30-conditions-bgp-conditions)# match-ext-
community-set excomm1
show running-config output
route-policy
  policy map1
    statement 30
      conditions
            bgp-conditions
              match-ext-community-set excomm1
```

Policy Result (Permit/Deny)

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 40
Switch(config-route-policy-policy-map1-stmt-40)# actions
Switch(config-route-policy-policy-map1-stmt-40-actions)# policy-result permit
show running-config output
route-policy
  policy map1
    statement 40
      actions
        policy-result permit
```

Set Community

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 50
Switch(config-route-policy-policy-map1-stmt-50)# actions
Switch(config-route-policy-policy-map1-stmt-50-actions)# bgp-actions
Switch(config-route-policy-policy-map1-stmt-50-actions-bgp-actions)# set-community add
comm1
show running-config output
route-policy
  policy map1
    statement 50
actions
        bgp-actions
          set-community add comm1
```

Set Ext-Community

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 60
Switch(config-route-policy-policy-map1-stmt-60)# actions
Switch(config-route-policy-policy-map1-stmt-60-actions)# bgp-actions
Switch(config-route-policy-policy-map1-stmt-60-actions-bgp-actions)# set-ext-community
add excomm1
show running-config output
route-policy
  policy map1
    statement 60
      actions
        bgp-actions
          set-ext-community add excomm1
```

Set Next Hop

```
Switch(config)# route-policy
Switch(config-route-policy)# policy map1
Switch(config-route-policy-policy-map1)# statement 70
Switch(config-route-policy-policy-map1-stmt-70)# actions
Switch(config-route-policy-policy-map1-stmt-70-actions)# bgp-actions
Switch(config-route-policy-policy-map1-stmt-70-actions-bgp-actions)# set-next-hop
10.1.1.1
show running-config output
route-policy
  policy map1
    statement 70
      actions
        bgp-actions
          set-next-hop 10.1.1.1
```

Configure Route Map on Extreme OS ONE Devices

You can configure a route map.

About This Task

Follow this procedure to configure a route map on Extreme OS ONE devices.

Procedure

1. Create a route map.

   To create a new route map with one or more rules:

   ```
   efa policy route-map create --name rmap_1
   --rule seq[5],action[permit]
   --rule seq[10],action[deny]
   ```

   This creates a route map named "rmap_1" with two rules:

   - Rule 5: Permit matching routes
   - Rule 10: Deny matching routes

2. Update a route map.

   Adds, updates, or removes device associations or updates route-map actions.

   ```
   efa policy route-map update
   --name rmap_1
   --operation add-device
   --ip 10.20.246.10-11

   efa policy route-map update
   --name rmap_1
   --operation remove-device
   --ip 10.20.246.10-11

   efa policy route-map update
   --name rmap_1
   ```

```
--rule seq[5],action[deny]
--operation update-action
```

Note

- **add-device** – Activate the route map on specified devices.
- **remove-device** – Remove the route map from specified devices.
- **update-action** – Change a rule's action from permit to deny (or vice versa).

3. Delete a route map.

   This removes routes (deletes specific rules or all rules) or deletes an entire route map. The following example removes rules 10 and 20 from the route map.

   ```
   efa policy route-map delete --name rmap_1 --seq 10,20
   ```

   The following example deletes all rules in a route map:

   ```
   efa policy route-map delete --name rmap_1 --seq all
   ```

4. View or list route maps.

   Displays configured route maps for managed devices. The following example displays the route maps for the specified devices:

   ```
   efa policy route-map list --ip 10.20.246.10-11
   ```

   The following example displays the detailed information of OOB entries:

   ```
   efa policy route-map list --ip 10.20.246.29 --detail
   ```

Configure Route Map Match Condition on Extreme OS ONE Devices

You can configure a route map match condition on Extreme OS ONE devices.

About This Task

Match conditions identify which routes the rule applies to. You can combine multiple conditions in a single rule.

Follow this procedure to configure a route map match condition on Extreme OS ONE devices.

Procedure

1. Match IPv4 or IPv6 prefixes.

   To match routes based on destination IP addresses:

   ```
   efa policy route-map-match create
    --name my_policy
    --rule seq[10],action[permit]
    --match-ipv4-prefix customers_in
   ```

   This applies rule 10 to routes whose destination addresses match the "customers_in" prefix list. Use --match-ipv6-prefix for IPv6 addresses.

2. Match BGP Communities.

   To match routes based on community values:

   ```
   efa policy route-map-match create
    --name my_policy
    --rule seq[20],action[permit]
    --match-community-list internal_routes
   ```

   This applies rule 20 to routes that belong to the "internal_routes" community.

3. Match Extended Communities.

   To match routes based on extended community values:

```
efa policy route-map-match create
 --name my_policy
 --rule seq[30],action[permit]
 --match-extcommunity-list vpn_routes
```

   This applies rule 30 to routes that match the "vpn_routes" extended community list.

4. Delete match criteria.

   Deletes match conditions from a route-map rule.

```
efa policy route-map-match delete
--name rmap_test
--rule "seq[10],action[permit]"
--match-ipv6-prefix prefix_1_in
```

Configure Route Map Set Action on Extreme OS ONE Devices

You can configure a route map set action on Extreme OS ONE devices.

About This Task

Actions define what happens to routes that match a rule. You can modify route attributes to influence routing decisions.

Follow this procedure to configure a route map set action on Extreme OS ONE devices.

Procedure

1. Modify the Next Hop.

   To change the next hop address for matching routes:

```
efa policy route-map-set create
 --name my_policy
 --rule seq[50],action[permit]
 --set-next-hop ipv4,192.168.1.1
```

   This directs matching routes through the gateway at 192.168.1.1. For IPv6, use --set-next-hop ipv6,2001:db8::1.

2. Set Community Values.

   To assign community values to matching routes:

```
efa policy route-map-set create
 --name my_policy
 --rule seq[60],action[permit]
 --set-community 65001:100,internet
```

   This marks matching routes with community values for easy identification and filtering.

3. Set Extended Communities.

   To assign extended community values for advanced route targeting:

```
efa policy route-map-set create
 --name my_policy
 --rule seq[70],action[permit]
 --set-extcommunity-rt 65001:300
```

4. Delete Set actions.

   Deletes route-map set actions.

```
efa policy route-map-set delete
 --name my_policy
 --rule seq[50],action[permit]
 --set-next-hop ipv4,192.168.1.1
```

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, 4.1.0*.

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, 4.1.0*.

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	
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
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, 4.1.0*.

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

Field	Identify Drift	Reconcile config	Idempotency	Comments
Large community attribute set criteria deleted	Yes	Yes	NA	Reconcile the set criteria for large community attribute
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	

Force Delete OOB Entries from Policy Configuration

You can delete out-of-band (OOB) entries from the XCO database and remove the configuration from associated devices.

About This Task

Follow this procedure to force delete OOB entries from policy configuration.

To delete OOB entries for every delete command on a policy type, you can use an optional command line parameter "--force". However, if the command line options contain XCO-managed entries, a warning message appears before you proceed to delete and unconfigure them from the associated devices.

Procedure

1. To delete all OOB rules associated with community-list “clist-oob” of “standard” type, run the **`efa policy community-list delete --name clist-oob --type standard --oob`** command.

   If the sequence numbers are specified and they include both XCO-managed and OOB entries, then all rules that match the provided sequence numbers associated with given community list will be deleted.

```
efa policy community-list list --type standard --ip 10.20.246.29,10.20.246.30
Name: clist-oob
Seq: 24
Action: deny
StdValue: 10:100 local-as
ExtValue:

Name: clist-oob
Seq: 23
Action: permit
StdValue: local-as
ExtValue:

IP Addresses:
+-----------+-----+--------------+-----------------+
|   Name    | Seq |  IP Address  |    App State    |
```

```
+-----------+-----+--------------+-----------------+
| clist-oob | 24  | 10.20.246.29 | cfg-in-sync     |
+-----------+-----+--------------+-----------------+
| clist-oob | 23  | 10.20.246.29 | cfg-not-managed |
+-----------+-----+--------------+-----------------+
```

2. To delete all OOB rules of the "extended" type, run the **efa policy extcommunity-list delete –name eclist-oob --type extended --seq all –-oob** command.

   You can also specify the sequence numbers you want to delete from the OOB list. The system will perform basic validation on the provided sequence numbers to ensure that they are associated with the correct community list.

```
efa policy extcommunity-list list --type extended --ip 10.20.246.29,10.20.246.30
Extended community list details:
Name: eclist-oob
Seq: 4
Action: deny
Route Target:
Site of Origin:
ExtValue: ^65000:.*_

Name: extcomm2
Seq: 4
Action: deny
Route Target:
Site of Origin:
ExtValue: ^65000:.*_

IP Addresses:
+------------------+-----+--------------+-----------------+
|       Name       | Seq |  IP Address  |    App State    |
+------------------+-----+--------------+-----------------+
| eclist-oob       | 4   | 10.20.246.29 | cfg-not-managed |
+------------------+-----+--------------+-----------------+
| extcomm2         | 4   | 10.20.246.29 | cfg-in-sync     |
+------------------+-----+--------------+-----------------+
```

3. To delete all OOB rules associated with the route map "rmap-oob", run the **efa policy route-map delete --name rmap-oob --seq all --oob** command.

```
efa policy route-map list --ip 10.20.246.29 --detail
Route Map details:
Name: rmap-oob
Seq: 32
Action: permit
Matches:
   match: matchExtcommunityList eclist-oob [cfg-not-managed]
   match: matchCommunityList clist-oob [cfg-not-managed]
   match: matchLargeCommunityList lclist-oob [cfg-not-managed]
Sets:
   set: setLargeCommunityList lclist-oob [cfg-not-managed]

Name: rmap1
Seq: 5
Action: permit
Matches:
   match: matchExtcommunityList extcommunitystd1
Sets:

Name: rmap1
Seq: 65535
Action: permit
Matches:
Sets:
```

```
   set: setCommunityValue 6550:125 internet local-as

IP Addresses:
+----------+-------+--------------+-----------------+
|   Name   |  Seq  |  IP Address  |    App State    |
+----------+-------+--------------+-----------------+
| rmap-oob | 32    | 10.20.246.29 | cfg-not-managed |
+----------+-------+--------------+-----------------+
| rmap1    | 5     | 10.20.246.29 | cfg-in-sync     |
+----------+-------+--------------+-----------------+
| rmap1    | 65535 | 10.20.246.29 | cfg-in-sync     |
+----------+-------+--------------+-----------------+
```

4. To delete all OOB rules associated with the prefix list "plist-oob", run the **`efa policy prefix-list delete --type ipv4 --name plist-oob --oob`** command.

```
efa policy prefix-list list --type ipv4 --ip 10.20.246.29

Prefix-list details:

Name: plist-oob
+------+---------+--------+------------+----+----+--------------+-----------------+
| Type | Seq num | Action |   Prefix   | Ge | Le |   DeviceIP   |    AppState     |
+------+---------+--------+------------+----+----+--------------+-----------------+
| ipv4 | 4       | permit | 1.1.1.1/32 |    |    | 10.20.246.29 | cfg-not-managed |
+------+---------+--------+------------+----+----+--------------+-----------------+
```

5. To delete all OOB rules associated with the prefix list "plis6t-oob" of IPv6 type, run the **`efa policy prefix-list delete --type ipv6 --name plist6-oob --oob`** command.

```
efa policy prefix-list list --type ipv6 --ip 10.20.246.29

Prefix-list details:

Name: plist6-oob
+------+---------+--------+---------------+----+----+--------------+-----------------+
| Type | Seq num | Action |    Prefix     | Ge | Le |   DeviceIP   |    AppState     |
+------+---------+--------+---------------+----+----+--------------+-----------------+
| ipv6 | 12      | deny   | 2006:db8::/44 |    |    | 10.20.246.29 | cfg-not-managed |
+------+---------+--------+---------------+----+----+--------------+-----------------+
```

6. To delete all OOB rules associated with the large community list "lclist-oob" of standard type, run the **`efa policy large-community-list delete --name lclist-oob --type standard --seq all --oob`** command.

```
efa policy large-community-list list --type standard --ip 10.20.246.29

Large Community list details:

Name: lclist-oob
Seq: 12
Action: permit
StdValue: 10:11:12
ExtValue:

Name: llist1
Seq: 5
Action: permit
StdValue: 10:20:30 50:60:70
ExtValue:

IP Addresses:
+------------+-----+--------------+-----------------+
|    Name    | Seq |  IP Address  |    App State    |
+------------+-----+--------------+-----------------+
```

```
| lclist-oob | 12  | 10.20.246.29 | cfg-not-managed |
+------------+-----+--------------+-----------------+
| llist1     | 5   | 10.20.246.29 | cfg-in-sync     |
+------------+-----+--------------+-----------------+
```

7. To delete all OOB rules associated with the QoS map "dscp-tc" of "dscp-tc-map" type, run the **efa policy qos map delete --name dscp-tc --type dscp-tc-map --oob** command.

```
 efa policy qos map list  --type dscp-tc-map --ip 10.20.246.29-30
QoS map details:
Name: dscp-tc
Type: dscp-tc-map
+------+----+----+
| DSCP | TC | DP |
+------+----+----+
| 0    | 0  | 2  |
+------+----+----+
| 1    | 0  | 2  |
+------+----+----+
| 2    | 0  | 2  |
+------+----+----+
| 3    | 0  | 2  |
+------+----+----+
Name: dscp-tc-oob
Type: dscp-tc-map
+------+----+----+
| DSCP | TC | DP |
+------+----+----+
| 0    | 2  | 2  |
+------+----+----+
| 1    | 2  | 2  |
+------+----+----+
| 2    | 2  | 2  |
+------+----+----+
| 4    | 2  | 2  |
+------+----+----+
| 5    | 2  | 2  |
+------+----+----+
Device Bindings:
+-------------+------+--------------+-----------------+
| dscp-tc     | 62   | 10.20.246.29 | cfg-in-sync     |
+-------------+------+--------------+-----------------+
| dscp-tc     | 62   | 10.20.246.30 | cfg-in-sync     |
+-------------+------+--------------+-----------------+
| dscp-tc     | 63   | 10.20.246.29 | cfg-in-sync     |
+-------------+------+--------------+-----------------+
| dscp-tc     | 63   | 10.20.246.30 | cfg-in-sync     |
+-------------+------+--------------+-----------------+
| dscp-tc-oob | 0    | 10.20.246.29 | cfg-not-managed |
+-------------+------+--------------+-----------------+
| dscp-tc-oob | 1    | 10.20.246.29 | cfg-not-managed |
```

8. To delete all OOB rules associated with the QoS map "pcp-tc" of "pcp-tc-map" type, run the **efa policy qos map delete --name pcp-tc --type pcp-tc-map --oob** command.

```
 efa policy qos map list  --type pcp-tc-map --ip 10.20.246.29-30

QoS map details:
Name: pcp-tc
Type: pcp-tc-map
+-----+----+----+
| PCP | TC | DP |
+-----+----+----+
```

```
| 0   | 5  | 1  |
+-----+----+----+
| 4   | 5  | 1  |
+-----+----+----+
| 5   | 5  | 1  |
+-----+----+----+
| 6   | 5  | 1  |
+-----+----+----+

Name: pcp-tc-oob
Type: pcp-tc-map
+-----+----+----+
| PCP | TC | DP |
+-----+----+----+
| 2   | 3  | 0  |
+-----+----+----+
| 3   | 4  | 0  |
+-----+----+----+
Device Bindings:
+------------+-----+--------------+-----------------+
| pcp-tc     | 6   | 10.20.246.29 | cfg-in-sync     |
+------------+-----+--------------+-----------------+
| pcp-tc     | 6   | 10.20.246.30 | cfg-in-sync     |
+------------+-----+--------------+-----------------+
| pcp-tc-oob | 2   | 10.20.246.29 | cfg-not-managed |
+------------+-----+--------------+-----------------+
| pcp-tc-oob | 3   | 10.20.246.29 | cfg-not-managed |
+------------+-----+--------------+-----------------+
```

9. To delete all OOB rules associated with the QoS map "tc-pcp" of "tc-pcp-map" type, run the **`efa policy qos map delete --name tc-pcp --type tc-pcp-map --oob`** command.

```
 efa policy qos map list  --type tc-pcp-map --ip 10.20.246.29-30
Name: tc-pcp
Type: tc-pcp-map
+----+----+-----+
| TC | DP | PCP |
+----+----+-----+
| 5  | 1  | 6   |
+----+----+-----+

Name: tc-pcp-oob
Type: tc-pcp-map
+----+----+-----+
| TC | DP | PCP |
+----+----+-----+
| 2  | 0  | 5   |
+----+----+-----+
| 3  | 1  | 4   |
+----+----+-----+

Device Bindings:
+------------+----+----+--------------+-----------------+
|    Name    | TC | DP |  IP Address  |    App State    |
+------------+----+----+--------------+-----------------+
| tc-pcp     | 5  | 1  | 10.20.246.29 | cfg-in-sync     |
+------------+----+----+--------------+-----------------+
| tc-pcp     | 5  | 1  | 10.20.246.30 | cfg-in-sync     |
+------------+----+----+--------------+-----------------+
| tc-pcp-oob | 2  | 0  | 10.20.246.29 | cfg-not-managed |
+------------+----+----+--------------+-----------------+
| tc-pcp-oob | 3  | 1  | 10.20.246.29 | cfg-not-managed |
+------------+----+----+--------------+-----------------+
```

10. To delete all OOB rules associated with the QoS service policy map "sp2", run the **`efa policy qos service-policy-map delete --name sp2 --oob`** command.

```
 efa policy qos service-policy-map list --ip 10.20.246.29-30

QoS Service Policy Map details:
+------+-----------------+----------------+---------+
| Name | Strict Priority |  DWRR Weights  |  Class  |
+------+-----------------+----------------+---------+
| sp2  | 4               | 25,25,25,25    | default |
+------+-----------------+----------------+---------+
| sp2  | 4               | 25,25,25,25    | default |
+------+-----------------+----------------+---------+
| sp5  | 3               | 20,20,20,20,20 | default |
+------+-----------------+----------------+---------+
| sp6  | 5               | 50,25,25       | default |
+------+-----------------+----------------+---------+

Device Bindings:
+--------------+-----------------+
|  IP Address  |    App State    |
+--------------+-----------------+
| 10.20.246.29 | cfg-in-sync     |
+--------------+-----------------+
| 10.20.246.30 | cfg-in-sync     |
+--------------+-----------------+
| 10.20.246.30 | cfg-not-managed |
+--------------+-----------------+
| 10.20.246.29 | cfg-not-managed |
+--------------+-----------------+
| 10.20.246.30 | cfg-not-managed |
+--------------+-----------------+
```

Example

```
efa policy community-list list --type standard --ip 10.20.246.29,10.20.246.30
Name: clist-oob
Seq: 24
Action: deny
StdValue: 10:100 local-as
ExtValue:

Name: clist-oob
Seq: 23
Action: permit
StdValue: local-as
ExtValue:

IP Addresses:
+-----------+-----+--------------+-----------------+
|   Name    | Seq |  IP Address  |    App State    |
+-----------+-----+--------------+-----------------+
| clist-oob | 24  | 10.20.246.29 | cfg-in-sync     |
+-----------+-----+--------------+-----------------+
| clist-oob | 23  | 10.20.246.29 | cfg-not-managed |
+-----------+-----+--------------+-----------------+

efa policy extcommunity-list delete -name eclist-oob --type extended --seq all --oob
```

To delete all OOB rules of the "extended" type, use the above command. You can specify the sequence numbers to be deleted from the OOB list. Basic validation is performed on the sequence numbers provided to ensure they are associated with the given community list.

```
efa policy extcommunity-list list --type extended --ip 10.20.246.29,10.20.246.30
Extended community list details:
Name: eclist-oob
Seq: 4
Action: deny
Route Target:
Site of Origin:
ExtValue: ^65000:.*_

Name: extcomm2
Seq: 4
Action: deny
Route Target:
Site of Origin:
ExtValue: ^65000:.*_

IP Addresses:
+------------------+-----+--------------+-----------------+
|       Name       | Seq |  IP Address  |    App State    |
+------------------+-----+--------------+-----------------+
| eclist-oob       | 4   | 10.20.246.29 | cfg-not-managed |
+------------------+-----+--------------+-----------------+
| extcomm2         | 4   | 10.20.246.29 | cfg-in-sync     |
+------------------+-----+--------------+-----------------+
efa policy route-map delete --name rmap-oob --seq all -force
 efa policy route-map list --ip 10.20.246.29 --detail
Route Map details:
Name: rmap-oob
Seq: 32
Action: permit
Matches:
   match: matchExtcommunityList eclist-oob [cfg-not-managed]
   match: matchCommunityList clist-oob [cfg-not-managed]
   match: matchLargeCommunityList lclist-oob [cfg-not-managed]
Sets:
   set: setLargeCommunityList lclist-oob [cfg-not-managed]

Name: rmap1
Seq: 5
Action: permit
Matches:
   match: matchExtcommunityList extcommunitystd1
Sets:

Name: rmap1
Seq: 65535
Action: permit
Matches:
Sets:
   set: setCommunityValue 6550:125 internet local-as

IP Addresses:
+----------+-------+--------------+-----------------+
|   Name   |  Seq  |  IP Address  |    App State    |
+----------+-------+--------------+-----------------+
| rmap-oob | 32    | 10.20.246.29 | cfg-not-managed |
+----------+-------+--------------+-----------------+
| rmap1    | 5     | 10.20.246.29 | cfg-in-sync     |
+----------+-------+--------------+-----------------+
| rmap1    | 65535 | 10.20.246.29 | cfg-in-sync     |
```

```
+----------+-------+--------------+-----------------+

efa policy prefix-list delete --type ipv4 --name plist-oob --oob
 efa policy prefix-list list --type ipv4 --ip 10.20.246.29

Prefix-list details:

Name: plist-oob
+------+---------+--------+------------+----+----+--------------+-----------------+
| Type | Seq num | Action |   Prefix   | Ge | Le |   DeviceIP   |    AppState     |
+------+---------+--------+------------+----+----+--------------+-----------------+
| ipv4 | 4       | permit | 1.1.1.1/32 |    |    | 10.20.246.29 | cfg-not-managed |
+------+---------+--------+------------+----+----+--------------+-----------------+

efa policy prefix-list delete --type ipv6 --name plist6-oob --oob
 efa policy prefix-list list --type ipv6 --ip 10.20.246.29

Prefix-list details:

Name: plist6-oob
+------+---------+--------+---------------+----+----+--------------+-----------------+
| Type | Seq num | Action |    Prefix     | Ge | Le |   DeviceIP   |    AppState     |
+------+---------+--------+---------------+----+----+--------------+-----------------+
| ipv6 | 12      | deny   | 2006:db8::/44 |    |    | 10.20.246.29 | cfg-not-managed |
+------+---------+--------+---------------+----+----+--------------+-----------------+

efa policy large-community-list delete --name lclist-oob --type standard --seq all --oob
 efa policy large-community-list list --type standard --ip 10.20.246.29

Large Community list details:

Name: lclist-oob
Seq: 12
Action: permit
StdValue: 10:11:12
ExtValue:

Name: llist1
Seq: 5
Action: permit
StdValue: 10:20:30 50:60:70
ExtValue:

IP Addresses:
+------------+-----+--------------+-----------------+
|    Name    | Seq |  IP Address  |    App State    |
+------------+-----+--------------+-----------------+
| lclist-oob | 12  | 10.20.246.29 | cfg-not-managed |
+------------+-----+--------------+-----------------+
| llist1     | 5   | 10.20.246.29 | cfg-in-sync     |
+------------+-----+--------------+-----------------+

efa policy qos map delete --name dscp-tc --type dscp-tc-map --oob
```

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 693)
- 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 693. 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 693. 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.

Policy Service QoS Support

Quality of Service (QoS) support includes setting QoS DSCP trust and defining & applying the QoS PCP and DSCP ingress and egress maps on Ethernet interfaces and Port Channels.

- The XCO-driven application of policy is dynamic and can vary depending on the port's role, whether it belongs to a fabric, tenant, port channel, or tenant endpoint group.

  

  Tip
  As a best practice, avoid running user-driven policy operations in parallel with fabric, tenant, port channel, and tenant endpoint group operations.

  To ensure that the fabric, tenant, port channel, and tenant endpoint group configurations are effective, run the **`show`** command before proceeding with the policy operations, and vice-versa.
- Before running the force operations, including deletion, ensure that you unbind the policies (QoS) from all the relevant targets (fabric, tenant, port, port channel, and tenant endpoint group) to avoid stale policies (QoS) in the system.
- Before running the QoS policy bind commands, remove the conflicting or additional OOB (Out Of Band) QoS configurations from the switches to ensure that the correct policies are applied to the ports.
- There is no support for a lossless hardware profile. Therefore, you must switch the configuration on SLX devices to a lossy hardware profile before provisioning QoS policies from XCO.
- There is no support for egress QoS maps. While XCO allows the configuration of egress QoS maps, as a best practice, do not configure any egress QoS maps from XCO due to limitations in SLX support of egress QoS maps.

QoS support in XCO 3.4.0 and later is as follows.

Product ID	SLX Version
Extreme 8720	20.4.3 or higher
Extreme 8520	20.4.3 or higher
SLX-9740	20.5.2 Partial support: Egress DSCP maps not supported.
Extreme 8820	20.5.2 Partial support: Egress DSCP maps not supported.

Product ID	SLX Version
SLX-9540	No support
SLX-9640	No support

Quality of Service (QoS) Implementation

There are two ways you can implement quality of service (QoS) in XCO:

- **QoS Map**

  QoS map defines mapping of header fields to QoS traffic classes on the device.

  The following header fields are considered for classification:

  - PCP field in the VLAN Tag
  - DSCP field in IP header

  You can map each of the header fields with the eight different traffic classes TC0-TC7. You can define QoS maps for mapping.

  For egress traffic, you can also specify the mapping between TC to PCP or DSCP field in the QoS maps. For example, you can define classification based on DSCP bits in the IP Header of the incoming packets. Each of the 64 DSCP values can be mapped to any one of the eight traffic classes.

- **QoS Profile**

  A QoS profile is a collection of information that can be applied on an interface or device that controls the queuing and dequeuing of traffic flow based on packet header information or internal queuing. QoS Profile defines the following:

  - Mapping of DSCP header from the IP header
  - CoS information (Priority Code Points) to internal traffic classes (TC)
  - Drop precedence (DP)
  - conversely mapping the TC to egress DSCP or PCP values

  The information includes global, interface, and flow configurations. The profile is a list of rules that define a desired QoS behavior.

  QoS profile is supported only for the physical interfaces. PO, VE, and breakout ports do not support QoS profile.

  You can use a QoS profile to define the classification and actions applied to the traffic such as the following:

  - Specify default traffic class for untagged frames.
  - Enable trust for DSCP values on all ingress traffic.
  - Specify maps for classification to traffic classes.
  - Specify maps for remarking rules (mutation maps).
  - Specify strict priority and/or DWRR weights.

- Specify interface-based rate-limiting or egress shaping.
- Specify flow-based rate-limiting or policing.

- **QoS Service Policy Map**

  QoS service policy maps define policies for scheduling, policing, and shaping network traffic.

  XCO 3.4.0 allows you to implement scheduling policies for better management of network traffic. With XCO 3.4.0, you can assign strict priority to certain traffic class queues while using DWRR (Deficit Weighted Round Robin) with customized weights for other traffic class queues.

  This enables you to efficiently manage network traffic and ensure that high-priority traffic is given the appropriate amount of bandwidth.

- **Association of Service Policy to QoS Profile**

  You can associate a QoS service policy map with a QoS profile by providing its name. Once associated, the policy settings (such as strict priority and DWRR) will be applied to all ports that the QoS profile is bound to.

Drift Reconcile and Idempotency

Configuration	Identify drift	Reconcile configuration	Idempotency
QoS Maps	Yes	Yes	No
QoS Profile	Yes	Yes	No

Configure QoS Map

You can use the `efa policy qos map` command to create, delete, update, and list QoS map.

About This Task

Follow this procedure to configure QoS map.

Procedure

1. To create a QoS map, run the **`efa policy qos map create`** command.

   For PCP and TC, the value list can have comma separated or range values between 0 to 7. For DSCP the value can be between 0 to 63. The following are the examples for PCP to TC and DSCP to TC mapping. Valid values for DP is 0 to 2. The **`efa policy qos map create`** command allows you to map many values to a single destination. The **`dscp-tc`** command allows you to map multiple dscp values to the same Traffic Class and Drop Precedence. For example, dscp[2-5;11] tc[2] dp[1] maps dscp values 2, 3, 4, 5, and 11 to traffic class with a drop precedence of 1.

   a. The following CLI creates a QoS map tenant1PcpToTCMap, which will classify packets with PCP values 0 to 7 in different Traffic Classes. You can provide values as a range:

   ```
   # efa policy qos map  create --name tenant1PcpToTCMap --type pcp-tc-map --rule
   "pcp[14-25],tc[7]" --rule "pcp[8],tc[1]" --rule "pcp[10],tc[1]"
   ```

   b. The following CLI creates a QoS map tenant1DscpToTCMap, which will classify packets with various DSCP values into different Traffic Classes. You can provide values as a range. Values that are not specified will be classified into default classes.

   ```
   # efa policy qos map create --name tenant1DscpToTCMap --type dscp-tc-map --rule
   "dscp[14-25],tc[7],dp[2]" --rule "dscp[8],tc[1]" --rule "dscp[10],tc[1],dp[1]"
   ```

2. To update QoS map, run the **`efa policy qos map update`** command.

   a. The following command updates the mapping for QoS map tenant1PcpToTCMap for the PCP values specified in the command. The original configuration prior to the “update” is as follows.

   ```
   # efa policy qos map  create --name tenant1PcpToTCMap --type pcp-tc-map --rule
   "pcp[14-25],tc[7]" --rule "pcp[8],tc[1]" --rule "pcp[10],tc[1]"
   ```

   Now map pcp 0 and pcp 1 to different traffic classes using the following command:

   ```
   % efa policy qos map update --name tenant1PcpToTCMap --type pcp-tc-map --rule
   “pcp[0],tc[6]” --rule “pcp[1],tc[7]”
   ```

   b. The following command updates the mapping for QoS map tenant1DscpToTCMap for the DSCP values specified in the command. Mapping that are not specified will remain unchanged. The original configuration prior to the “update” is as follows.

   ```
   % efa policy qos map create --name tenant1DscpToTCMap --type dscp-tc-map --rule
   "dscp[14-25],tc[7],dp[2]" --rule "dscp[8],tc[1]" --rule "dscp[10],tc[1],dp[1]"
   ```

   Now provide additional DSCP mappings beyond those originally specified using the following command:

   ```
   % efa policy qos map update --name tenant1DscpToTCMap --type dscp-tc-map --rule
   “dscp[0-7],tc[1]” --rule “dscp[22-38],tc[5],dp[1]”
   ```

3. To delete QoS map, run the **`efa policy qos map delete`** command.

   The **`efa policy qos map delete`** command deletes the configuration from XCO database. Deletion of QoS map fails if it is part of a profile.

   a. The following command deletes the QoS map tenant1PcpToTCMap from XCO database:
   ```
   % efa policy qos map delete --name tenant1PcpToTCMap --type pcp-tc-map
   ```

   b. The following command deletes the user defined DSCP to TC mapping for DSCP values 0 to 7 in the QoS map tenant1DscpToTCMap. This will cause the traffic with DSCP values 0 to 7 to be classified based on the system defaults.
   ```
   % efa policy qos map delete --name tenant1DscpToTCMap --type dscp-tc-map -rule
   "dscp[0-7]"
   ```

4. To list QoS maps, run the **`efa policy qos map list`** command.

   When you provide the `--ip` option, the app-state of configuration will be displayed for a given device.

   Note
   The app-state only conveys the state of the map on the device and not the bindings to the interfaces. To see the app-state of the maps applied on an interface, check the output of QoS profile list commands.

   a. The following command lists configuration details of the QoS Map tenant1DscpToTCMap10 that was created with parameters:
   ```
   % efa policy qos map create --name tenant1DscpToTCMap10 --type dscp-tc-map --rule
   "dscp[0-7],tc[3],dp[1]" --rule "dscp[8],tc[1]" --rule "dscp[10],tc[1],dp[1]"

   % efa policy qos map list --name tenant1DscpToTCMap10 --type dscp-tc-map

   QoS Map details:
   Name: tenant1DscpToTCMap10
   +-------+-------+-------+
   |  DSCP |   TC  |   DP  |
   +-------+-------+-------+
   |   0   |   3   |   1   |
   +-------+-------+-------+
   |   1   |   3   |   1   |
   +-------+-------+-------+
   |   2   |   3   |   1   |
   +-------+-------+-------+
   |   3   |   3   |   1   |
   +-------+-------+-------+
   |   4   |   3   |   1   |
   +-------+-------+-------+
   |   5   |   3   |   1   |
   +-------+-------+-------+
   |   6   |   3   |   1   |
   +-------+-------+-------+
   |   7   |   3   |   1   |
   +-------+-------+-------+
   |   8   |   1   |   0   |
   +-------+-------+-------+
   |  10   |   1   |   1   |
   +-------+-------+-------+
   ```

   b. The following command lists configuration details of the QoS map tenant1DscpToTCMap with the app state on the specified devices:
   ```
   % efa policy qos map list --name tenant1DscpToTCMap10 --type dscp-tc-map --ip
   10.20.245.1-2
   ```

```
QoS Map details:
Name: tenant1DscpToTCMap10
+-------+-------+-------+
|  DSCP |   TC  |   DP  |
+-------+-------+-------+
|   0   |   3   |   1   |
+-------+-------+-------+
|   1   |   3   |   1   |
+-------+-------+-------+
|   2   |   3   |   1   |
+-------+-------+-------+
|   3   |   3   |   1   |
+-------+-------+-------+
|   4   |   3   |   1   |
+-------+-------+-------+
|   5   |   3   |   1   |
+-------+-------+-------+
|   6   |   3   |   1   |
+-------+-------+-------+
|   7   |   3   |   1   |
+-------+-------+-------+
|   8   |   1   |   0   |
+-------+-------+-------+
|  10   |   1   |   1   |
+-------+-------+-------+


Device Bindings:
+------------------------+------+----------------+----------------+
|           Name         | DSCP |   IP Address   |   App State    |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   0  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   1  |   10.20.245.1  |  cfg-refreshed |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   2  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   3  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   4  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   5  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   6  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   7  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |   8  |   10.20.245.1  |  cfg-refreshed |
+------------------------+------+----------------+----------------+
| tenant1DscpToTCMap10   |  10  |   10.20.245.1  |   cfg-in-sync  |
+------------------------+------+----------------+----------------+
```

c. The following command lists configuration details of all the QoS maps of type dscp-tc-map with the app state on the specified device:

```
% efa policy qos map list --type dscp-tc-map --ip 10.20.245.1
QoS Map details:
Name: tenant1DscpToTCMap15
+--------+-------+-------+
|  DSCP  |   TC  |   DP  |
+--------+-------+-------+
|   10   |   3   |   1   |
+--------+-------+-------+
|   15   |   3   |   1   |
+--------+-------+-------+
```

```
Name: oobDscpToTCMap1
+--------+-------+-------+
|  DSCP  |   TC  |   DP  |
+--------+-------+-------+
|   20   |   5   |   1   |
+--------+-------+-------+
|   25   |   5   |   1   |
+--------+-------+-------+


Device Bindings:
+-----------------------+-------+----------------+----------------=-+
|           Name        |  DSCP |   IP Address   |    App State     |
+-----------------------+-------+----------------+------------------+
| tenant1DscpToTCMap15  |   10  |   10.20.245.1  |    cfg-in-sync   |
+-----------------------+-------+----------------+------------------+
| tenant1DscpToTCMap15  |   15  |   10.20.245.1  |   cfg-refreshed  |
+-----------------------+-------+----------------+------------------+
|   oobDscpToTCMap1     |   20  |   10.20.245.1  |  cfg-not-managed |
+-----------------------+-------+----------------+------------------+
|   oobDscpToTCMap1     |   25  |   10.20.245.1  |  cfg-not-managed |
+-----------------------+-------+----------------+------------------+
```

Note
For more information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Drop Precedence Support for Extreme OS ONE

Drop precedence support enables Extreme OS ONE devices to classify packets based on congestion priority. During congestion, packets with a higher drop precedence value are discarded before packets with a lower drop precedence value. This capability provides consistent QoS behavior across SLX and Extreme OS ONE platforms

The QoS drop precedence implementation supports the following capabilities:

- Configure drop precedence values in QoS traffic-class and remark maps
- Create, update, and manage QoS mappings by using a unified EFA CLI
- Maintain consistent QoS behavior across Extreme OS ONE and SLX platforms
- Define packet handling priority during congestion scenarios

The supported drop precedence values are:

Drop Precedence Value	Description
0	Low drop precedence — least likely to be dropped
1	Medium drop precedence
2	High drop precedence — most likely to be dropped

QoS Map Types

The following QoS map types support drop precedence:

QoS Map Type	Description
`dscp-tc-map`	Maps DSCP values to traffic class values
`pcp-tc-map`	Maps PCP values to traffic class values
`tc-dscp-map`	Maps traffic class values to DSCP values
`tc-pcp-map`	Maps traffic class values to PCP values

Note

On Extreme OS ONE devices:

- `dscp-tc-map` and `pcp-tc-map` must use the same map name.
- `tc-dscp-map` and `tc-pcp-map` must use the same map name.

Validation and Error Handling

The valid drop precedence range is: 0..2

- Example Invalid Configuration

```
efa policy qos map create
--type tc-pcp-map
--name tcMap
--rule "pcp[7],tc[2],dp[10]"
```

- Error Output

```
Error : Invalid dp value in (10). The valid value is 0..2
```

Device CLI Configuration Example

The following example shows the equivalent configuration directly on the E1OS device CLI.

```
kvm-x86-64# conf
kvm-x86-64(config)# qos

kvm-x86-64(config-qos)# traffic-class-map tcMap
kvm-x86-64(config-qos-traffic-class-map-tcMap)# dscp 5 traffic-class 4 drop-precedence 1
kvm-x86-64(config-qos-traffic-class-map-tcMap)# cos 4 traffic-class 3 drop-precedence 0

kvm-x86-64(config-qos)# remark-map reMap
kvm-x86-64(config-qos-remark-map-reMap)# traffic-class 4 drop-precedence 2 cos 3
kvm-x86-64(config-qos-remark-map-reMap)# traffic-class 3 drop-precedence 1 dscp 2
```

Configure QoS Maps with Drop Precedence on Extreme OS ONE

You can configure QoS maps with drop precedence on Extreme OS ONE.

Before You Begin

Before you configure QoS drop precedence mappings, ensure that:

- XCO is deployed and operational
- The target Extreme OS ONE device is onboarded in XCO

- QoS policies are enabled on the device
- You have administrator privileges

About This Task

Follow this procedure to configure QoS maps with drop precedence on Extreme OS ONE.

Procedure

1. Create a DSCP-to-Traffic-Class Map.

```
efa policy qos map create
--type dscp-tc-map
--rule "dscp[4],tc[3],dp[1]"
--name tcMap

+--------------+---------------+------+-----+----+----+
| QoS Map Name | Map Type      | Dscp | Pcp | Tc | Dp |
+--------------+---------------+------+-----+----+----+
| tcMap        | dscp-tc-map   | 4    | 0   | 3  | 1  |
+--------------+---------------+------+-----+----+----+

QoS Map details
Success
```

2. Create a PCP-to-Traffic-Class Map.

```
efa policy qos map create
--type pcp-tc-map
--name tcMap
--rule "pcp[1],tc[5],dp[2]"
```

3. Create a Traffic-Class-to-DSCP Map.

```
efa policy qos map create
--type tc-dscp-map
--name reMap
--rule "dscp[1],tc[5],dp[0]"
```

4. Create a Traffic-Class-to-PCP Map.

```
efa policy qos map create
--type tc-pcp-map
--name tcMap
--rule "pcp[7],tc[2],dp[1]"
```

5. Update QoS Maps.

   To update the DSCP-to-traffic-class mapping, run the following command:

```
efa policy qos map update
--type dscp-tc-map
--name tcMap
--rule "dscp[4],tc[7],dp[2]"

+--------------+---------------+------+-----+----+----+
| QoS Map Name | Map Type      | Dscp | Pcp | Tc | Dp |
+--------------+---------------+------+-----+----+----+
| tcMap        | dscp-tc-map   | 4    | 0   | 7  | 2  |
+--------------+---------------+------+-----+----+----+

QoS Map details
Success
```

6. Verify the running configuration.

   To verify the QoS configuration, run the **`show running-config qos`** command. The output displays all QoS maps and their drop precedence settings.

```
qos
  traffic-class-map tcMap
    dscp 5 traffic-class 4 drop-precedence 1
    cos 4 traffic-class 3 drop-precedence 0
  !
  remark-map reMap
    traffic-class 4 drop-precedence 2 cos 3
    traffic-class 3 drop-precedence 1 dscp 2
  !
!
```

Example

`

- **Example 1: Voice over IP (VoIP) Protection**

  Scenario: Your organization uses VoIP for critical communications and wants to ensure these calls are never dropped during congestion.

  Solution: Assign VoIP traffic (typically DSCP EF = 46) a low drop precedence (0) to protect it during congestion.

```
efa policy qos map create --type dscp-tc-map --name voip-protection --rule
"dscp[46],tc[7],dp[0]"
```

- **Example 2: Tiered Service Levels**

  Scenario: You want to support three service tiers: Premium, Standard, and Best-Effort.

  Solution: Create a map that assigns different drop precedence levels based on customer class:

  - Premium traffic: Drop Precedence 0
  - Standard traffic: Drop Precedence 1
  - Best-Effort traffic: Drop Precedence 2

```
efa policy qos map create --type dscp-tc-map --name tiered-service --rule
"dscp[10],tc[1],dp[0];dscp[20],tc[3],dp[1];dscp[0],tc[0],dp[2]"
```

Configure QoS Profile

You can use the `efa policy qos profile` command to create, delete, update, bind, unbind, and list QoS profile.

About This Task

Follow this procedure to configure QoS profile on a fabric.

Procedure

1. Create the maps.
   a. Use maps for the association of DSCP or PCP values to TC and the converse mapping. Maps associate a header value or defined TC to another value. The map types

referenced by a QoS Profile are pcp-tc-map, dscp-tc-map, tc-pcp-map, and tc-dscp-map. The first step is to define the maps to which the profile will be referenced.

Run the following command to create one of each map. The command for creating the four map types are similar and are of the form “from value” -> “to value”. For example, pcp-to-tc type takes a “pcp” value and maps it to a “TC + DP” combination.

```
efa policy qos map create --name mapName  --type pcp-to-tc --rule
“pcp[0],tc[7],dp[1]” --rule “pcp[1],tc[6]” --rule “pcp[2-5]tc[5],dp[2]”

efa policy qos map create --name tc2pcp --type tc-pcp-map --rule "pcp[5],tc[2]"
--rule "pcp[2],tc[3],dp[1]"

efa policy qos map create --name dscp2tc --type dscp-tc-map --rule "dscp[5],tc[2]"
--rule "dscp[2],tc[3],dp[1]"

efa policy qos map create --name tc2dscp --type tc-dscp-map --rule
"dscp[63],tc[0],dp[1]" --rule "dscp[0],tc[1],dp[0]"
```

b. Once a “rule” is added to a map to change the behavior of the rule, use the “update” version of the command. For example, one of the pcp-to-tc rule maps pcp 0 to TC 7 DP1. If it is now desired to map pcp 0 to TC 6 DP 1, run the following command:

```
efa policy qos map update --name mapName  --type pcp-to-tc --rule
“pcp[0],tc[6],dp[1]”
```

c. If the desired mapping is no longer desired and to set the mapping of pcp 0 back to the default setting, delete the existing rule by running the following command:

```
efa policy qos map delete --name mapName  --type pcp-to-tc --rule “pcp[0]”
```

2. Once the maps are created, create a QoS Profile that reference to the existing maps. Run the following command to create a profile:

```
efa policy qos profile create --name qf1 --trust dscp --pcp-tc pcp2tc --dscp-tc
dscp2tc --tc-pcp tc2pcp --tc-dscp tc2dscp
```

This command contains maps to be used on an interface for different mappings or mutations and if manipulation of DSCP on Layer 2 interfaces are allowed. For more details on how the profile is applied to a device, see step 4.

3. Run the following command to apply the QoS profile to the fabric:

QoS Profiles are applied at the fabric level and will be applied to all interfaces of all devices that belong to the specified fabric.

```
efa policy qos profile bind --name qf1 --fabric MyFabric
```

4. To delete QoS profile after it is applied, run the following unbind command:

The delete command of QoS profile removes (or defaults) specified parameter from an existing profile. If no parameter is specified, the QoS profile will be deleted. This change will be made to the XCO database. It might take effect on devices if the configuration is applied on the devices.

```
efa policy qos profile unbind --name qf1 --fabric MyFabric
```

- The following CLI removes the specific rules, for example, the DSCP to TC map associated with the profile:

```
% efa policy qos profile delete --name fabricProfile1 --dscp-tc
```

- The following CLI deletes the entire profile, fabricProfile1:

```
% efa policy qos profile delete --name fabricProfile1
```

5. To list QoS profiles, run the **`efa policy qos profile list`** command.

   If you provide an IP option, the app state of the configuration will be displayed for the given devices.

   - The following CLI lists the configuration details of QoS profile qosProfile1:

```
% efa policy qos profile list --name qosProfile1

QoS Profile details:
Name: qosProfile1
Pcp->TC: Map1
```

   - The following CLI lists configuration details of QoS profile qosProfile10 where multiple user defined maps and trust are specified as 'auto':

```
% efa policy qos profile list --name qosProfile10

QoS Profile details:
Name: qosProfile10
Trust: auto
Pcp->TC: Map1
Dscp->TC: Map2
TC->Dscp: Map3
```

   - The following CLI lists configuration details of QoS profile qosProfile10 when you provide a range of devices (Binding or app states of given QoS profile on multiple devices):

```
% efa policy qos profile list --name qosProfile10 --ip 10.20.245.1-3

QoS Profile details:
Name: qosProfile10
Trust: auto
Pcp->TC: Map1
Dscp->TC: Map2
TC->Dscp: Map3

Device Bindings:
+--------------+---------------+--------------+
|  Name        |  IP Address   |   App State  |
+--------------+---------------+--------------+
| qosProfile10 | 10.20.245.1   |  cfg-in-sync |
+--------------+---------------+--------------+
|              | 10.20.245.2   |  cfg-in-sync |
+--------------+---------------+--------------+
|              | 10.20.245.3   |  cfg-in-sync |
+--------------+---------------+--------------+
```

   - The following CLI lists configuration details of QoS profile qosProfile10 when you provide a device (Binding or app states of given QoS profile on a single device). This command displays the app state of QoS Profile on all interfaces of the device:

```
% efa policy qos profile list --name qosProfile10 --ip 10.20.245.1

QoS Profile details:
Name: qosProfile10
Trust: auto
Pcp->TC: Map1
Dscp->TC: Map2
TC->Dscp: Map3
Device Bindings:
+--------------+---------------+--------------+--------------+
|  Name        |  IP Address   |   Interface  |   App State  |
+--------------+---------------+--------------+--------------+
| qosProfile10 | 10.20.245.1   | Ethernet 0/1 |  cfg-in-sync |
+--------------+---------------+--------------+--------------+
|              |               | Ethernet 0/2 |  cfg-in-sync |
```

```
+--------------+---------------+--------------+--------------+
|              |               | Ethernet 0/3 |  cfg-in-sync |
+--------------+---------------+--------------+--------------+
|              |               | Ethernet 0/4 |  cfg-in-sync |
+--------------+---------------+--------------+--------------+
```

- The following CLI lists configuration details of QoS profile qosProfile10 when you provide a device and interface. This command displays the app state of the QoS profile parameters on a given interface of the specified device:

```
% efa policy qos profile list --name qosProfile10 --ip 10.20.245.1 --interface
"Ethernet 0/3"

QoS Profile details:
Name: qosProfile10
Trust: auto
Pcp->TC: Map1
Dscp->TC: Map2
TC->Dscp: Map3

Device Bindings:
+--------------+---------------+--------------+----------+--------------+
|  Name        |  IP Address   |  Interface   | Parameter|   App State  |
+--------------+---------------+--------------+----------+--------------+
| qosProfile10 | 10.20.245.1   | Ethernet 0/3 |  Trust   |  cfg-in-sync |
+--------------+---------------+--------------+----------+--------------+
|              |               |              |  Pcp->TC |  cfg-in-sync |
+--------------+---------------+--------------+----------+--------------+
|              |               |              | Dscp->TC |  cfg-in-sync |
+--------------+---------------+--------------+----------+--------------+
|              |               |              | TC->Dscp |  cfg-in-sync |
+--------------+---------------+--------------+----------+--------------+
```

6. Complete the following configuration on SLX devices:

```
efa policy qos map create --type dscp-tc --name dscp2tc --rule "dscp[1],tc[1],dp[1]"
efa policy qos map create --type pcp-tc --name pcp2tc --rule "pcp[2],tc[5],dp[2]"
efa policy qos map create --type tc-dscp --name tc2dscp --rule "tc[4],dp[2],dscp[44]"
efa policy qos map create --type tc-pcp --name tc2pcp --rule "tc[3],dp[1],pcp[3]"
efa policy qos profile create --name qosProfile1 --rule --dscp-tc[dscp2tc] --tc-
dscp[tc2dscp]  --pcp-tc[pcp2tc] --tc-pcp[tc2pcp] --trust[dscp]
efa policy qos profile bind --profile qosProfile1 --fabric fabric1
```

The following example describes how an SLX configuration is applied on devices for QoS profile:

a. The required QoS maps are created on the fabric devices.

On devices of fabric1	On devices of fabric2
qos map traffic-class-dscp tc2Dscp map traffic-class 4 drop-precedence 2 to dscp 44 ! qos map dscp-traffic-class dscp2tc map dscp 1 to traffic-class 1 drop-precedence 1 !	SLX# show running-config qos qos map cos-traffic-class pcp2tc map cos 2 to traffic-class 5 drop-precedence 2 ! qos map traffic-class-cos tc2pcp map traffic-class 3 drop-precedence 1 to cos 3 !

b. The `trust dscp` configuration is applied to all the L2 interfaces.

L2 interface in access mode	L2 interface in trunk mode
SLX# show running-config interface Ethernet 0/27 interface Ethernet 0/27 switchport switchport mode access switchport access vlan 1 qos trust dscp no shutdown !	interface Ethernet 0/25 switchport switchport mode trunk switchport trunk tag native-vlan qos trust dscp no shutdown !

c. PCP to TC and TC to PCP maps are applied on all L2 interfaces in trunk mode.

```
interface Ethernet 0/25
switchport
switchport mode trunk
switchport trunk tag native-vlan
qos trust dscp
qos traffic-class-cos tc2pcp
qos cos-traffic-class pcp2tc
no shutdown
!
```

When you set the trust to auto in the profile, the configuration is as follows.

```
interface Ethernet 0/25
switchport
switchport mode trunk
switchport trunk tag native-vlan
qos traffic-class-cos tc2pcp
qos cos-traffic-class pcp2tc
no shutdown
!
```

d. DSCP to TC is applied to all the L3 interfaces and L2 interfaces if qos trust is set to DSCP.

- L2 interface in trunk mode

```
interface Ethernet 0/25
switchport
switchport mode trunk
switchport trunk tag native-vlan
qos trust dscp
qos traffic-class-cos tc2pcp
qos cos-traffic-class pcp2tc
qos dscp-traffic-class dscp2tc
```

```
no shutdown
!
```

- L2 interface in access mode

```
interface Ethernet 0/27
switchport
switchport mode access
switchport access vlan 1
qos trust dscp
qos dscp-traffic-class dscp2tc
no shutdown
!
```

- L3 interface

```
interface Ethernet 0/26
qos dscp-traffic-class dscp2tc
no shutdown
!
```

e. TC to DSCP maps are applied to all the L3 interfaces and L2 interfaces irrespective of `qos trust dscp` setting.

- L2 interface in trunk mode

```
interface Ethernet 0/25
switchport
switchport mode trunk
switchport trunk tag native-vlan
qos trust dscp
qos traffic-class-cos tc2pcp
qos cos-traffic-class pcp2tc
qos traffic-class-dscp tc2Dscp
qos dscp-traffic-class dscp2tc
qos remark dscp
no shutdown
!
```

- L2 interface in access mode (trust enabled)

```
interface Ethernet 0/27
switchport
switchport mode access
switchport access vlan 1
qos trust dscp
qos traffic-class-dscp tc2Dscp
qos dscp-traffic-class dscp2tc
qos remark dscp
no shutdown
!
```

- L2 interface in access mode (trust not enabled)

```
interface Ethernet 0/28
switchport
switchport mode access
switchport access vlan 1
qos traffic-class-dscp tc2Dscp
qos remark dscp
no shutdown
!
```

- L3 interface

```
interface Ethernet 0/26
qos traffic-class-dscp tc2Dscp
qos dscp-traffic-class dscp2tc
no shutdown
!
```

Application of QoS Profile

You can apply QoS profiles to the interfaces in a hierarchical manner.

You can apply a QoS profile at the following three conceptual levels:

1. **Fabric Interface level**: Configuration is applied to the physical interfaces that are part of devices in a fabric.
2. **Tenant interface level**: Configuration is applied to all the interfaces that are part of the tenant.
3. **Individual Interface level**: Configuration is applied to a subset of interfaces within a specified Tenant.

Note
XCO 3.3.0 and later supports application of profile only at fabric level.

The QoS profiles applied at individual interface level will have highest precedence than the QoS profiles applied at the fabric level. If no profile is applied on an interface, system level defaults will take effect.

Device Application

A QoS profile contains five configuration elements but not all the five configuration elements are applied to every interface. How the information is applied to an interface is dependent on the "switchport mode" of the interface. The following table describes the mapping between QoS configuration element and switchport mode:

	Switchport mode "l2 access"	Switchport mode "l2 trunk"	Switchport mode "l3"
Trust DSCP	Applied	Applied	N/A
dscp-tc	Applied if "Trust = DSCP"	Applied if "Trust = DSCP"	Applied
tc-dscp	Applied	Applied	Applied
pcp-tc	N/A	Applied	N/A
tc-pcp	N/A	Applied	N/A

L3 mode

DSCP is always trusted and the "Trust DSCP" element is ignored or not required to associate dscp-tc or tc-dscp. L3 interfaces do not operate on L2 headers and therefore the pcp-tc and tc-pcp have no bearings on the L3 interface. The results of applying the five configuration on an L3 port will result in the following configuration on a device:

```
interface Ethernet 0/5
 qos traffic-class-dscp tc2dscp
 qos dscp-traffic-class dscp2tc
 no shutdown
!
```

L2 Access mode

No mapping or mutation is performed using PCP on access interfaces to the defined maps are not applied. The dscp-tc map is only applied to the interface if the "trust" element is

set to “dscp”. If the Trust element is not specified or is set explicitly to “auto” the dscp-tc mapping will not be configured on the interface. If tc-dscp is contained in the profile it will be configured on the interface along with the "remark dscp” configuration. The results of applying the five configuration on an L2 Access port will result in the following configuration on a device:

```
interface Ethernet 0/1
 switchport
 switchport mode access
 switchport access vlan 1
 qos trust dscp
 qos remark dscp
 qos traffic-class-dscp tc2dscp
 qos dscp-traffic-class dscp2tc
 no shutdown
!
```

L2 Trunk mode

L2 Trunks use the PCP of the l2 header to map to and from different traffic classes on the device so any pcp-tc or tc-pcp specified configuration is mapped to an interface operating in this mode. As for the use of dscp-tc or tc-dscp are handled as is for l2 access interface. The dscp-tc map is only applied to the interface if the “trust” element is set to “dscp”. If the Trust element is not specified or is set explicitly to “auto” the dscp-tc mapping will not be configured on the interface. If tc-dscp is contained in the profile it will be configured on the interface along with the "remark dscp” configuration. The results of applying the five configuration on an L2 Trunk port would result in the following configuration on a device:

```
interface Ethernet 0/2
 switchport
 switchport mode trunk
 switchport trunk tag native-vlan
 qos trust dscp
 qos remark dscp
 qos traffic-class-cos tc2pcp
 qos cos-traffic-class pcp2tc
 qos traffic-class-dscp tc2dscp
 qos dscp-traffic-class dscp2tc
 no shutdown
```

Tenant and Port Level Binding

You can apply profiles at both the tenant and port levels.

The following diagram illustrates how concurrent profiles are applied at different levels in a hierarchy and how they affect system configuration. The QoS profiles applied to individual interfaces have the highest priority, while those applied to the fabric level have the lowest priority. If no profile is applied on an interface, the default settings at the system level will be used.

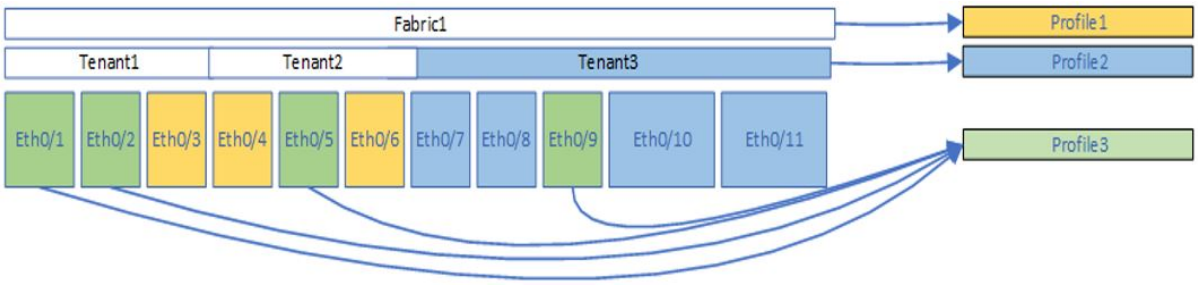

Use the following command to bind or unbind a profile at fabric, tenant or a subset of interfaces within a tenant:

```
% efa policy qos profile [bind | unbind] --profile <qosProfileName> --fabric|tenant
<fabric-name|teannt-name>
```

QoS Profile Application to Fabric

```
% efa policy qos profile [bind | unbind] --profile <qosProfileName>
```

`--fabric <fabricName>` indicates the specified profile is being applied to a set of interfaces belonging to the specified fabric.

Use the QoS profile command to:

- Apply the QoS Profile settings on the fabric devices and interfaces which is considered as a default fabric configuration. Tenant or interface configuration will override this configuration.
- Remove the QoS Profile settings from all fabric devices and interfaces. Only device level configuration with fabric configuration will be removed. Tenant or interface-level configurations will not be affected.

Example:

1. The following command applies the configuration of a QoS Profile named fabricProfile1 to all devices and/or interfaces in “fabric1”:
   ```
   % efa policy qos profile bind --profile fabricProfile1 --fabric fabric1
   ```
2. The following command removes the configuration of a QoS profile named fabricProfile1 from all devices and/or interfaces on fabric1, on which the configuration has been applied:
   ```
   % efa policy qos profile unbind --profile fabricProfile1 --fabric fabric1
   ```

QoS Profile Application to Tenant

```
% efa policy qos profile [bind | unbind] --profile <qosProfileName>
```

- `--tenant <tenantName>` = This indicates the specified profile is being applied to a set of interfaces belonging to the specified tenant.
- `--port “IP[ifName,ifName ...], IP[ifName, …]”` - This parameter is optional and only used to associate a policy to a subset of interfaces within a tenant.
- `IP[ifName,ifName...]` - This is an identifier of the device and interface on which the specified policy is to be applied.
- `ifName` - This is an abbreviated name of standard interface types. For example, `eth = ethernet`. The `ifName` can contain a comma separated list of interfaces on the device. For example, 1.1.1.1[0/1,0/2].

Use the QoS profile command to:

- Apply the QoS profile settings on devices and interfaces assigned to a tenant. Note that this configuration will overwrite any fabric level QoS configuration. However, if a member interface of a tenant already has a QoS profile assigned at the “individual interface” level, this command will not overwrite or change the existing configuration.
- Remove the QoS profile settings from all devices and interfaces of a Tenant. When removing a tenant QoS binding, if there is an existing profile binding on the fabric which contains tenant, then the tenant QoS configuration is removed and the QoS configuration specified in the fabric binding is applied.
- Apply the commands to all interfaces within a tenant by removing the `– interface` option or to a subset of interfaces within the tenant by providing a list of the interfaces using the – interface option.
- Ensure that first you create a tenant and then apply the profile to the interfaces within the tenant because application of QoS profiles to an interface requires the interfaces to be a member of subset within a tenant.

Example of bind operation:

1. The following command applies the configuration of a QoS profile named tenantProfile1 to all devices and/or interfaces in “tenant1”:
   ```
   % efa policy qos profile bind --profile tenantProfile1 --tenant tenant1
   ```
2. The following example applies the configuration of QoS profile "fabricProfile1" to all devices and interfaces in fabric "fabric1".
   ```
   efa policy qos profile bind --profile fabricProfile1 --fabric fabric1
   ```
3. The following example applies the configuration of QoS profile "fabricProfile2" to all internal ports in fabric "fabric1", but excludes applying the configuration on any edge interfaces. These ports were previously bound with "fabricProfile1" in the above example.
   ```
   efa policy qos profile bind --profile fabricProfile2 --fabric fabric1 --port “fabric-internal”
   ```
4. The following example applies the configuration of a QoS profile named "tenantProfile1" to all devices or interfaces in "tenent1".
   ```
   efa policy qos profile bind --profile tenantProfile1 --tenant tenant1
   ```

Example fo unbind operation:

1. The following example removes the configuration of a specified QoS profile from all devices and interfaces on a specified fabric.
   ```
   efa policy qos profile unbind --profile fabricProfile1 --fabric fabric1
   ```
2. The following example removes the configuration of a QoS profile named "tenantProfile2" from all devices or interfaces on "tenant1" on which the configuration has been applied.
   ```
   efa policy qos profile unbind --profile tenantProfile1 --tenant tenant1
   ```
3. The following example removes the configuration of a QoS profile named "tenantProfile2" from device IP 1.1.1.1 ethernet 0/1 (ethernet 0/1 is a member of tenant1 and po3, po4)
   ```
   efa policy qos profile unbind --profile tenantProfile2 --tenant tenant1 --port “1.1.1.1[0/1]” --po po3,po4
   ```

Bind Operation Example of Concurrently Applied Profiles

This topic describes examples on configuring fabric, tenants and QoS profiles and illustrates how the binding of QoS profiles within the binding hierarchy operates.

Example:
To start the base configuration, a Fabric1 is created that contains Ethernet interface Eth0/1-Eth0/11. There are also three QoS profiles created, Profile1-Profile3 that have not yet been bound at any level in the hierarchy.

1. Three Tenants are created. Tenant1 has eth0/1-0/3, Tenant2 has eth0/4-eth0/6, and Tenant3 has eth0/7-eth0/11. Profile1 is applied to Fabric1:
   ```
   efa policy qos profile bind --profile Profile1 --fabric fabric1
   ```
2. Apply Profile2 to Tenant 3:
   ```
   efa policy qos profile bind --profile Profile2 --tenant tenant3
   ```
3. Apply Profile3 directly to few discrete interfaces. To achieve the configuration depicted in the following diagram, run three “bind” command. One command for each subset of interfaces within the three Tenants.
   ```
   %efa policy qos profile bind --profile Profile3 --tenant tenant1 --
   port “1.1.1.1[0/1,0/2]”
   %efa policy qos profile bind --profile Profile3 --tenant tenant2 --
   port “1.1.1.1[0/5]”
   %efa policy qos profile bind --profile Profile3 --tenant tenant3 --
   port “1.1.1.1[0/9]"--po "po3,po4”
   ```

Example:
To start the base configuration, a Fabric1 is created that contains Ethernet interface Eth0/1-Eth0/11. There are also three QoS profiles created, Profile1-Profile3 that have not yet been bound at any level in the hierarchy. After the configuration, the only allowable options are either to bind a profile to Fabric1 or create tenants and apply profiles to those tenants.

1. Create three tenants and assign Profile2 to Tenant3:
   ```
   efa policy qos profile bind --profile Profile2 --tenant tenant3
   ```
2. Apply Profile3 directly to few discrete interfaces. To achieve the configuration depicted in the following diagram the user will need to issue 3 “bind” command. One command for each subset of interfaces within the 3 Tenants:
   ```
   %efa policy qos profile bind --profile Profile3 --tenant tenant1 --
   port “1.1.1.1[0/1,0/2]”
   %efa policy qos profile bind --profile Profile3 --tenant tenant2 --
   port “1.1.1.1[0/5]”
   %efa policy qos profile bind --profile Profile3 --tenant tenant3 --
   port “1.1.1.1[0/9]"--po "po3,po4”
   ```
3. Apply Profile1 to Fabric1.
   ```
   % efa policy qos profile bind --profile Profile1 --fabric fabric1
   ```

Unbind Operation Example of Concurrently Applied Profiles

The following is the initial configuration for the two example is as follows.

```
% efa policy qos profile bind --profile Profile1 --fabric fabric1
% efa policy qos profile bind --profile Profile2 --tenant tenant3
% efa policy qos profile bind --profile Profile3 --tenant tenant1 --port
“1.1.1.1[0/1,0/2]”
% efa policy qos profile bind --profile Profile3 --tenant tenant2 --port “1.1.1.1[0/5]”
% efa policy qos profile bind --profile Profile3 --tenant tenant3 --port “1.1.1.1[0/9]"--
po "po3,po4”
```

Example:

1. Use the following command to remove the binding of Profile3 from eth 0/9:
   ```
   efa policy qos profile unbind --profile Profile3 --tenant tenant3 --port "1.1.1.1[0/9]"
   ```
   The configuration of "Profile3" is removed from Eth0/9. Eth0/9 is reconfigured to contain the QoS configuration contained in "Profile2" as eth0/9 is a member of Tenant3, and Tenant 3 has Profile2 actively applied.
2. Use the following command to remove the binding of Profile2 from tenant3:
   ```
   efa policy qos profile unbind --profile Profile2 --tenant tenant3
   ```
   The configuration of "Profile2" is removed from Eth0/7, Eth0/8, Eth0/10, and Eth0/11. These four interfaces will be reconfigured to contain the QoS configuration contained in "Profile1".

Note

Since Profile3 is applied to Eth0/9 using the –-interface option when removing the Tenant3 binding even though Eth0/9 is member of Tenant3, the configuration of Eth0/9 is unchanged and still contains the configuration associated with Profile3.

Apply QoS Profile

You can bind or unbind a profile using the **`efa policy qos profile [bind | unbind] --profile <qosProfileName> --fabric|tenant --<fabric-name|teannt-name>`**.

About This Task

Follow this procedure to bind or unbind a profile at fabric, tenant or a subset of interfaces within a tenant.

Use the **`efa policy qos profile [bind | unbind] --profile <qosProfileName> --fabric|tenant --<fabric-name|teannt-name>`** command to apply or remove the QoS profile.

- You can apply the QoS profile settings on devices and interfaces of a fabric. This configuration is treated as the default fabric configuration and any tenant or interface configuration will override this configuration.
- Remove the QoS profile settings from all devices and interfaces of a fabric. It removes only those configuration on devices which has fabric configuration. Any tenant or interface level configurations will not be affected.

Note

For more information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

Run the following command to bind or unbind the QoS profile configuration:

```
% efa policy qos profile [bind | unbind] --profile <qosProfileName> --fabric|tenant
<fabric-name|teannt-name>
```

Example

1. The following CLI applies the configuration of a QoS profile fabricProfile1 to all devices and interfaces in fabric1:
```
% efa policy qos profile bind --profile fabricProfile1 --fabric fabric1
```
2. The following CLI removes the configuration of a QoS profile fabricProfile1 from all devices and interfaces in fabric1:
```
% efa policy qos profile unbind --profile fabricProfile1 --fabric fabric1
```

QoS Profile Binding Enhancement

.Use the information in the following topics to learn about the improved binding of Quality of Service (QoS) profiles.

Introduction

1. XCO 3.5.0 and later has enhanced the binding of Quality of Service (QoS) profiles. It supports incremental application and the reconciliation of failed bindings caused by conflicting configurations present on devices. It includes out-of-band (OOB) changes to interface mode settings, moving interfaces into and out of port-channels, concurrent execution of tenant CRUD (Create, Read, Update, Delete) operations related to Policy Objects (POs), and Endpoint Groups (EPGs) while QoS profile binding is in progress, and more. The following cases provide some common scenarios and explain how error handling is performed.
2. If a device in the fabric cannot be reached due to connectivity issues, no commands will be sent to any device. You must wait for the connectivity to be restored before attempting the command again.
3. In case the error is legitimate, such as an incorrect switch-port mode on an interface or QoS on LAG member that cannot be applied, the binding process may partially succeed on some devices. However, the failed devices will be reconciled, and the issue will be corrected. The app-state (cfg-refreshed) will be displayed for all failed interfaces, and it will be updated to (cfg-in-sync) after the configuration is reconciled.
4. Successful devices will have a result of SUCCESS while failed devices will have a result of PARTIAL. There is no "Rollback" in a partial success scenario.
5. When a device binding is in a "Partial" state, it can be automatically applied when relevant events are processed, and the failed bindings are attempted for reconciliation. However, there is a possibility that the policy service may never receive any subsequent events to correct the failed bindings. In such cases, the user will need to re-issue the failed bind command to complete the provisioning of failed bindings.

User can check if recovery is pending by listing interfaces in “Partial” state.

```
efa policy qos profile list --ip 10.20.246.1
Name: profileTen
Trust:  dscp
Dscp->Tc:  dscp2tcTen
Service Policy Map: spTen, out
Name: profileTenIF
Trust:  dscp
Dscp->Tc:  dscp2tcTenIF
Service Policy Map: spTen, out
IP Addresses:
+--------------+-------------+-----------+-----------------+-----------+---------------+
```

```
| profileTen   | 10.20.246.1 | eth 0/29  | tenant           | tenantAB  | cfg-in-sync   |
+--------------+-------------+-----------+------------------+-----------+---------------+
| profileTen   | 10.20.246.1 | eth 0/30  | tenant           | tenantAB  | cfg-in-sync   |
+--------------+-------------+-----------+------------------+-----------+---------------+
| profileTenIF | 10.20.246.1 | eth 0/11  | tenant-interface | tenantAB  | cfg-refreshed |
+--------------+-------------+-----------+------------------+-----------+---------------+
After the recovery is complete, app-state of interfaces will be changed from “cfg-refreshed” to “cfg-
in-sync".
 efa policy qos profile bind --name profileTen --tenant tb
Error :  Device 10.64.208.16 not reachable. Please retry after verifying the inputs and connectivity
issues

 efa policy qos profile bind --name profileFab --fabric fabric1
+------------------+-----------+-----------+
| Qos Profile Name | Bind Type | Bind Name |
+------------------+-----------+-----------+
| profileFab       | fabric    | fabric1   |
+------------------+-----------+-----------+

Qos Profile details
+--------------+---------+-------------------------------------------------------------------+-----------------+
|  IP Address  | Result  |                            Reason                                 | Rollback reason |
+--------------+---------+-------------------------------------------------------------------+-----------------+
| 10.20.246.30 | Success |                                                                   |                 |
+--------------+---------+-------------------------------------------------------------------+-----------------+
| 10.20.246.29 | Partial | IFs: [0/15 0/16 0/24 0/29 0/30 0/1 0/13 0/10|                     |                 |
|              |         | Reason: netconf rpc [error] QoS config on LAG                     |                 |
|              |         | member is not allowed'                                            |                 |
+--------------+---------+-------------------------------------------------------------------+-----------------+


 efa policy qos profile bind --name profileFabIF --fabric fabric1 --port fabric-internal
+------------------+-----------+-----------+
| Qos Profile Name | Bind Type | Bind Name |
+------------------+-----------+-----------+
| profileFabIF     | fabric    | fabric1   |
+------------------+-----------+-----------+

Qos Profile details

+--------------+---------+-------------------------------------------------------------------
+----------------+
|  IP Address  | Result  |                            Reason
| Rollback reason| |
+--------------+---------+-------------------------------------------------------------------
+----------------+
| 10.20.246.29 | Partial | IFs: [64] Err: Reason:
netconf rpc [error] '%Error:          |               |
|              |         | Policy-Map not found'
|                |
+--------------+---------+-------------------------------------------------------------------
+----------------+
| 10.20.246.30 | Partial | IFs: [64] Err: Reason:
netconf rpc [error] '%Error:          |               |
|              |         | Policy-Map not found'
|                |
+--------------+---------+-------------------------------------------------------------------
+----------------+
```

Sometimes, when multiple actions are happening at the same time (concurrent execution of tenant CRUD operations and QoS profile binding by policy service), there can be issues. For instance, if a tenant service is modifying a switchport mode on an interface by creating an EPG while a QoS profile (tenantIF) is being bound, it can cause the tenantIF profile binding

to fail. This happens because the switchport mode change could occur before the binding process is complete. If this occurs, the operation will result in failure, and the device will display the following output:

```
efa policy qos profile bind --name profileTenIF --tenant tenantAB --port
"10.20.246.1[0/11],10.20.246.2[0/11]"

+------------------+-----------+-----------+
| Qos Profile Name | Bind Type | Bind Name |
+------------------+-----------+-----------+
| profileTenIF     | tenant    | tenantAB  |
+------------------+-----------+-----------+

Qos Profile details
+-------------+---------+-------------------------------------------------------------------
+-----------------+
| IP Address  | Result  |                            Reason
| Rollback reason |
+-------------+---------+-------------------------------------------------------------------
+-----------------+
| 10.20.246.2 | Partial | IFs: [0/11] Err: Reason:
netconf rpc [error] '%Error: Enable |                  |
|             |         | QoS trust DSCP'
|                 |
+-------------+---------+-------------------------------------------------------------------
+-----------------+
| 10.20.246.1 | Partial | IFs: [0/11] Err: Reason:
netconf rpc [error] '%Error: Enable |                  |
|             |         | QoS trust DSCP'
|                 |
+-------------+---------+-------------------------------------------------------------------
+-----------------+
```

To ensure successful provisioning, re-issue the binding after the switchport mode changes have been processed, following any switchport mode changes made to the EPG. Alternatively, if the tenantIF binding is applied immediately after the creation of the EPG, it is assumed that the switchport mode changes due to EPG creation have been processed before the application of binding, resulting in successful provisioning.

Policy Service CLI Changes

Option `-po` uses PO name string instead of PO ID string for the **`efa policy qos bind/ unbind`** command.

Following are the examples of an existing command and a modified command:

- Existing Command:
  ```
  efa policy qos profile bind --name profile2 --tenant tenant1 --po "1,3"
  efa policy qos profile unbind --name profile2 --tenant tenant1 --po "1,3"
  ```
- Modified Command:
  ```
  efa policy qos profile bind --name profile2 --tenant tenant1 --po "po1,po3"
  efa policy qos profile unbind --name profile2 --tenant tenant1 --po "po1,po3"
  ```

Configure QoS Service Policy Map

You can configure a QoS service policy map.

About This Task

Follow this procedure to configure a QoS service policy map.

For more information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

1. Run the following command to create a QoS service policy map:

   The command creates the QoS service policy map configuration in the XCO database. This configuration is not pushed to the device. Configuration will only be pushed when the service policy map is associated with a QoS profile and the QoS profile is bound.

```
% efa policy qos service-policy-map create ?
  --name string        Name of the QoS service policy map
  --rule stringArray   Rule in format
  "strict-priority[3],dwrr[25;25;25;25;0],class[default]"

Values of strict-priority as defined below
  <0  No strict priority queue; all are DWRR
1  Traffic Class 7 strict priority queue; rest are DWRR
2  Traffic Class 6 through 7 strict priority queues; rest are DWRR
3  Traffic Class 5 through 7 strict priority queues; rest are DWRR
4  Traffic Class 4 through 7 strict priority queues; rest are DWRR
5  Traffic Class 3 through 7 strict priority queues; rest are DWRR
6  Traffic Class 2 through 7 strict priority queues; rest are DWRR
7  Traffic Class 1 through 7 strict priority queues; rest are DWRR >
```

   - For weights, the value list can have semicolon separated weights for traffic class queues which do not have strict priority. The sum of all the weights must equal to 100.
   - SLXOS Default: All TC queues are strict priority. TC7 has the highest priority and TC0 has the lowest.

   The following command creates a QoS service policy map named servicePolicy1, which specifies strict priority for traffic classes 5-7 and specifies DWRR weights for traffic classes 0-4:

```
% efa policy qos service-policy-map create --name servicePolicy1 --rule "strict-
priority[3],dwrr[25;25;25;25;0],class[default]"
```

2. Run the following command to update a QoS service policy map:

   Note

   You cannot update the content of a policy map when an actively bound QoS profile references the map.

```
% efa policy qos service-policy-map update ?
  --name string        Name of the QoS service policy map
  --rule stringArray   Rule in format
  "strict-priority[3],dwrr[25;25;25;25;0],class[default]"

Values of strict-priority as defined below
  <0  No strict priority queue
1  Traffic Class 7 strict priority queue
2  Traffic Class 6 through 7 strict priority queues
3  Traffic Class 5 through 7 strict priority queues
```

```
4  Traffic Class 4 through 7 strict priority queues
5  Traffic Class 3 through 7 strict priority queues
6  Traffic Class 2 through 7 strict priority queues
7  Traffic Class 1 through 7 strict priority queues>
Only “default” for class attribute is supported
```

The value list can have semicolon separated weights for traffic class queues which do not have strict priority. The sum of all the weights must equal to 100.

The following command updates the QoS service policy map named servicePolicy1 for strict priority value of 2 and new weights:

```
% efa policy qos service-policy-map update --name servicePolicy1 --rule “strict-
priority[2],dwrr[50;25;25;0;0;0],class[default]”
```

Running the **`efa policy qos service-policy-map update`** command updates the configuration on the switch as follows.

```
S4(config-policymap-class)# do show running-config policy-map
policy-map servicePolicy1
 class default
  scheduler strict-priority 2 dwrr 50 25 25 0 0 0
```

3. Run the following command to delete a QoS service policy map:

```
% efa policy qos service-policy-map delete?
  --name string        Name of the QoS service policy map
```

The following command deletes the QoS service policy map servicePolicy1:

```
% efa policy qos service-policy-map delete --name servicePolicy1
```

4. Run the following command to list a QoS service policy map:

```
% efa policy qos service-policy-map list?
  --name string        Name of the QoS service policy map
  --ip   string      IP Address
```

- The following command lists configuration details of the QoS service policy map servicePolicy1:

```
% efa policy qos service-policy-map list --name servicePolicy1

QoS Service Policy Map details:
Name: servicePolicy1
+------------------+-------------------+---------+
|  Strict Priority |   DWRR weights    | Class   |
+------------------+-------------------+---------+
|         2        | 50,25,25,0,0,0    | default |
+------------------+-------------------+---------+
```

- The following command lists configuration details of the QoS map servicePolicy1 along with the app state on specified devices:

```
% efa policy qos service-policy-map list --name servicePolicy1 --ip 10.20.245.1-2
QoS Service Policy Map details:
Name: servicePolicy1
+------------------+-------------------+---------+
|  Strict Priority |   DWRR weights    | Class   |
+------------------+-------------------+---------+
|         2        | 50,25,25,0,0,0    | default |
+------------------+-------------------+---------+

Device Bindings:
+----------------+----------------+
|   IP Address   |   App State    |
+----------------+----------------+
|   10.20.245.1  |   cfg-in-sync  |
```

```
+----------------+----------------+
|   10.20.245.2  |   cfg-in-sync  |
+----------------+----------------+
```

XCO QoS Support for Extreme OS ONE

QoS maps in Extreme OS ONE define the mapping of header fields (PCP and DSCP) to traffic classes (TC0-TC7). Users can create QoS maps for ingress and egress traffic, specifying mappings between TC and PCP/DSCP fields.

QoS Profile

A QoS profile is a named logical construct that specifies classification and actions for traffic. It can be used to:

- Set default traffic class for untagged frames
- Enable trust for DSCP values on ingress traffic
- Set maps for classification to traffic classes.
- Set maps for remarking rules (mutation maps)
- Configure strict priority and DWRR weights
- Set interface-based rate-limiting or egress shaping
- Specify flow-based rate-limiting or policing

A QoS profile contains all necessary QoS settings for devices, covering global, interface, and flow configurations. It's essentially a rule-based definition of desired QoS behavior

CLI Interfaces for QoS Configuration

The XCO CLI interface, options, and values for QoS maps and profiles remain largely consistent for Extreme OS ONE, with a few key differences.

Extreme OS ONE Configuration Notes

- **Map Name Consistency**: When configuring TC to DSCP/PCP and DSCP/PCP to TC maps for a specific fabric, tenant, or interface, the map names must be identical as described in the following table:

Map Type	SLX	Extreme OS ONE
PCP-TC	map1	map1
TC-PCP	map2	map2
DSCP-TC	map3	map1
PCP-TC	map4	map2

- **Trust Configuration**: By default, no traffic is considered trusted. Even with QoS trust set to auto on L2 ports, explicit trust configuration (for example, trust PCP) is required on Extreme OS ONE devices.
- **Remark Maps**: On Extreme OS ONE devices, PCP-TC and DSCP-TC maps are treated as remark maps.

- **Drop Precedence**: It is not supported on Extreme OS ONE. As a best practice, avoid configuring this value to prevent drift during DRC.
- **QoS Profile Priority**: The priority order remains the same on the device:
  - Tenant interface (highest)
  - Tenant level
  - Fabric level (lowest)

QoS Configuration Examples

Use this topic to learn about the examples of Quality of Service (QoS) configuration using the Command-Line Interface (CLI).

QoS Map Configuration

```
efa policy qos map create --name Tcmap --type pcp-tc-map --rule pcp[0],tc[7] --rule
pcp[1],tc[6] --rule pcp[2-3],tc[5]
efa policy qos map create --name Tcmap --type dscp-tc-map --rule dscp[14-25],tc[7] --rule
dscp[8],tc[1] --rule dscp[10],tc[1]
efa policy qos map create --name Remap --type tc-pcp-map --rule pcp[5],tc[2] --rule
pcp[2],tc[3]
efa policy qos map create --name Remap --type tc-dscp-map --rule dscp[63],tc[0] --rule
dscp[0],tc[1]
efa policy qos map list --ip 10.10.10.75
```

QoS Map Configuration Details

```
Name: Tcmap
Type: pcp-tc-map
+-----+----+----+
| PCP | TC | DP |
+-----+----+----+
| 0   | 6  | 0  |
+-----+----+----+
| 1   | 5  | 0  |
+-----+----+----+
| 2   | 4  | 0  |
+-----+----+----+
| 3   | 4  | 0  |
+-----+----+----+

Device Bindings:
+-------+-----+--------------+-------------+
| Name  | PCP |  IP Address  |  App State  |
+-------+-----+--------------+-------------+
| Tcmap | 0   | 10.10.10.75 | cfg-in-sync |
+-------+-----+--------------+-------------+
| Tcmap | 1   | 10.10.10.75 | cfg-in-sync |
+-------+-----+--------------+-------------+
| Tcmap | 2   | 10.10.10.75 | cfg-in-sync |
+-------+-----+--------------+-------------+
| Tcmap | 3   | 10.10.10.75 | cfg-in-sync |
+-------+-----+--------------+-------------+

Name: Tcmap
Type: dscp-tc-map
+------+----+----+
| DSCP | TC | DP |
+------+----+----+
| 8    | 1  | 0  |
```

```
+------+----+----+
| 10   | 1  | 0  |
+------+----+----+
| 14   | 7  | 0  |
+------+----+----+
| 15   | 7  | 0  |
+------+----+----+
| 16   | 7  | 0  |
+------+----+----+
| 17   | 7  | 0  |
+------+----+----+
| 18   | 7  | 0  |
+------+----+----+
| 19   | 7  | 0  |
+------+----+----+
| 20   | 7  | 0  |
+------+----+----+
| 21   | 7  | 0  |
+------+----+----+
| 22   | 7  | 0  |
+------+----+----+
| 23   | 7  | 0  |
+------+----+----+
| 24   | 7  | 0  |
+------+----+----+
| 25   | 7  | 0  |
+------+----+----+

Device Bindings:
+-------+------+--------------+-------------+
| Name  | DSCP |  IP Address  |  App State  |
+-------+------+--------------+-------------+
| Tcmap | 8    | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 10   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 14   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 15   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 16   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 17   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 18   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 19   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 20   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 21   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 22   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 23   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 24   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+
| Tcmap | 25   | 10.10.10.75 | cfg-in-sync |
+-------+------+--------------+-------------+

Name: Remap
Type: tc-pcp-map
+----+----+-----+
```

```
| TC | DP | PCP |
+----+----+-----+
| 2  | 0  | 5   |
+----+----+-----+
| 3  | 0  | 2   |
+----+----+-----+

Device Bindings:
+-------+----+----+--------------+-------------+
| Name  | TC | DP |  IP Address  |  App State  |
+-------+----+----+--------------+-------------+
| Remap | 2  | 0  | 10.10.10.75 | cfg-in-sync |
+-------+----+----+--------------+-------------+
| Remap | 3  | 0  | 10.10.10.75 | cfg-in-sync |
+-------+----+----+--------------+-------------+

Name: Remap
Type: tc-dscp-map
+----+----+------+
| TC | DP | DSCP |
+----+----+------+
| 0  | 0  | 63   |
+----+----+------+
| 1  | 0  | 0    |
+----+----+------+

Device Bindings:
+-------+----+----+--------------+-------------+
| Name  | TC | DP |  IP Address  |  App State  |
+-------+----+----+--------------+-------------+
| Remap | 0  | 0  | 10.10.10.75 | cfg-in-sync |
+-------+----+----+--------------+-------------+
| Remap | 1  | 0  | 10.10.10.75 | cfg-in-sync |
+-------+----+----+--------------+-------------+
--- Time Elapsed: 370.679119ms ---
```

QoS Profile Configuration

```
efa policy qos profile create --name pf1 --pcp-tc Tcmap --dscp-tc Tcmap --tc-pcp Remap
--tc-dscp Remap --trust dscp
```

QoS Profile Binding

```
efa policy qos profile bind --name pf1 --fabric fabric2

efa policy qos profile list --ip 10.10.10.75

QoS profile details:

Name: pf1
Trust:  dscp
Pcp->Tc:  Tcmap
Dscp->Tc:  Tcmap
Tc->Pcp:  Remap
Tc->Dscp: Remap

IP Addresses:
+------+--------------+-----------+-----------+-----------+-------------+
| Name |  IP Address  | Interface | Bind Type | Bind Name |  App State  |
+------+--------------+-----------+-----------+-----------+-------------+
| pf1  | 10.10.10.75 | eth 0/1   | fabric    | fabric2   | cfg-in-sync |
+------+--------------+-----------+-----------+-----------+-------------+
| pf1  | 10.10.10.75 | eth 0/2   | fabric    | fabric2   | cfg-in-sync |
+------+--------------+-----------+-----------+-----------+-------------+
```

```
show running-config qos traffic-class-map
qos
  traffic-class-map Tcmap
    cos 3 traffic-class 4
    cos 0 traffic-class 6
    cos 1 traffic-class 5
    cos 2 traffic-class 4
    dscp 8 traffic-class 1
    dscp 10 traffic-class 1
    dscp 24 traffic-class 7
    dscp 17 traffic-class 7
    dscp 18 traffic-class 7
    dscp 22 traffic-class 7
    dscp 14 traffic-class 7
    dscp 16 traffic-class 7
    dscp 21 traffic-class 7
    dscp 23 traffic-class 7
    dscp 19 traffic-class 7
    dscp 20 traffic-class 7
    dscp 25 traffic-class 7
    dscp 15 traffic-class 7
  !
!
show running-config qos remark-map
qos
  remark-map Remap
    traffic-class 3 cos 2
    traffic-class 2 cos 5
    traffic-class 0 dscp 63
    traffic-class 1 dscp 0
  !
!
show running-config qos interface
qos
  interface ethernet 0/1
    input
      trust dscp
      traffic-class-map Tcmap
    !
    output
      remark dscp
      remark-map Remap
    !
  !
```

Policy Debug Drift and Reconcile

Before running the **`efa policy debug drift-reconcile`** command, ensure the **`efa inventory drift-reconcile execute`** command has fully completed.

QoS Maps and Profile Update

To update QoS profiles or maps, you must first unbind them from fabrics or tenants. Only "to" values can be updated in QoS maps, such as:

- TC in PCP-to-TC maps
- PCP in TC-to-PCP maps

To update QoS maps and profiles, complete the following steps:

1. Unbind QoS profile from fabric or tenant.

2. Update QoS map or profile.
3. Rebind if necessary.

```
efa policy qos map update --name Tcmap --type pcp-tc-map --rule pcp[0],tc[6] --rule
pcp[1],tc[5] --rule pcp[2-3],tc[4]
+--------------+------------+------+-----+----+----+
| QoS Map Name |  Map Type  | Dscp | Pcp | Tc | Dp |
+--------------+------------+------+-----+----+----+
| Tcmap        | pcp-tc-map | 0    | 3   | 4  | 0  |
+--------------+------------+------+-----+----+----+
| Tcmap        | pcp-tc-map | 0    | 0   | 6  | 0  |
+--------------+------------+------+-----+----+----+
| Tcmap        | pcp-tc-map | 0    | 1   | 5  | 0  |
+--------------+------------+------+-----+----+----+
| Tcmap        | pcp-tc-map | 0    | 2   | 4  | 0  |
+--------------+------------+------+-----+----+----+
QoS Map details
Success
--- Time Elapsed: 217.104705ms ---

efa policy qos profile update --name pf1 --pcp-tc Tcmap --dscp-tc Tcmap --tc-pcp Remap
--tc-dscp Remap --trust dscp
+------------------+-------+--------+---------+--------+---------+----------------
+-----------+
| QoS Profile Name | Trust | Pcp-Tc | Dscp-Tc | Tc-
Pcp | Tc-Dscp | Service-Policy | Direction |
+------------------+-------+--------+---------+--------+---------+----------------
+-----------+
| pf1              | dscp  | Tcmap  | Tcmap   | Remap
| Remap   |                |           |
+------------------+-------+--------+---------+--------+---------+----------------
+-----------+
QoS Profile details
Success
```

Qos Profile Unbinding

```
efa policy qos profile unbind --name pf1 --fabric fabric2
```

QoS Profile Deletion

```
efa policy qos profile delete --name pf1
```

QoS Map Deletion

```
efa policy qos map delete --name Tcmap --type=pcp-tc-map
efa policy qos map delete --name Tcmap --type=dscp-tc-map
efa policy qos map delete --name Remap --type=tc-pcp-map
efa policy qos map delete --name Remap --type=tc-dscp-map
```

Policy Service Support on Extreme OS ONE

Use this topic to learn about the equivalent SLX OS and Extreme OS ONE running configurations for QoS maps, remark maps, service policies, and QoS profiles across the fabric, fabric interface, tenant, and tenant interface.:

Policy configuration on Extreme OS ONE is fully supported without any changes to the existing XCO CLI workflow. All SLX policy-related CLI commands continue to function on Extreme OS ONE in XCO 4.0.2.

The following capabilities are supported:

1. XCO 4.0.2 supports **QoS Maps** and **QoS Profiles** for Extreme OS ONE.
2. QoS map types **DSCP→TC**, **TC→DSCP**, **PCP→TC**, and **TC→PCP** are supported on Extreme OS ONE. *Drop Precedence parameters are supported only where applicable and not in all mappings for this release.*
3. QoS Profiles can be applied at the **Fabric (non-clause)** level, **Tenant**, and **Tenant Interface** levels on Extreme OS ONE.

Fabric-Level QoS Profile and Configuration (SLX OS vs OS ONE)

Use this topic to learn about the QoS profile configuration at the Fabric level, including the creation of QoS maps, service policy maps, and the associated QoS profile through XCO. It also compares the resulting SLX OS and OS ONE running configurations.

QoS Map and Profile Creation (XCO CLI)

- The following maps are created to define PCP–TC, DSCP–TC, TC–PCP, and TC–DSCP translations:

```
efa policy qos map create --name pcp2tcFab --type pcp-tc-map --rule "pcp[5-7],tc[2]"

efa policy qos map create --name dscp2tcFab --type dscp-tc-map --rule "dscp[5],tc[2]"
--rule "dscp[2],tc[3],dp[1]"

efa policy qos map create --name tc2pcpFab --type tc-pcp-map --rule "pcp[5],tc[2]"
--rule "pcp[2],tc[3],dp[1]"

efa policy qos map create --name tc2dscpFab --type tc-dscp-map --rule
"dscp[63],tc[0],dp[1]" --rule "dscp[0],tc[1],dp[0]"
```

- A service-policy map is defined as follows:

```
efa policy qos service-policy-map create --name servicepolicy1 --rule "strict-
priority[0],dwrr[2;2;3;4;5;6;0;78],class[default]"
```

- The QoS profile is then created, referencing the previously defined maps:

```
efa policy qos profile create --name profileFab --trust dscp --pcp-tc pcp2tcFab
--dscp-tc dscp2tcFab --tc-dscp tc2dscpFab --tc-pcp tc2pcpFab --service-policy
name[servicepolicy1],dir[out]
```

QoS Profile Details

Attribute	Value
Profile Name	profileFab
Trust Mode	dscp
PCP → TC Map	pcp2tcFab
DSCP → TC Map	dscp2tcFab
TC → PCP Map	tc2pcpFab
TC → DSCP Map	tc2dscpFab
Service Policy Map	servicepolicy1 (direction: out)

SLX OS Running Configuration

- QoS Maps

```
qos map cos-traffic-class pcp2tcFab
 map cos 5 to traffic-class 2 drop-precedence 0
 map cos 6 to traffic-class 2 drop-precedence 0
 map cos 7 to traffic-class 2 drop-precedence 0
!

qos map traffic-class-cos tc2pcpFab
 map traffic-class 2 drop-precedence 0 to cos 5
 map traffic-class 3 drop-precedence 1 to cos 2
!

qos map traffic-class-dscp tc2dscpFab
 map traffic-class 0 drop-precedence 1 to dscp 63
 map traffic-class 1 drop-precedence 0 to dscp 0
!

qos map dscp-traffic-class dscp2tcFab
 map dscp 2 to traffic-class 3 drop-precedence 1
 map dscp 5 to traffic-class 2 drop-precedence 0
!
```

- Service Policy

```
policy-map servicepolicy1
  class default
    scheduler strict-priority 0 dwrr 2 2 3 4 5 6 0 78
  !
!
```

- Interface Configuration

```
interface Ethernet 0/10
 service-policy out servicepolicy1
 qos traffic-class-dscp tc2dscpFab
 qos dscp-traffic-class dscp2tcFab
 no shutdown
!
```

OS ONE Running Configuration

- QoS Maps and Remark Maps

```
qos
  traffic-class-map pcp2tcFab
    cos 7 traffic-class 2
    cos 5 traffic-class 2
    cos 6 traffic-class 2
  !

traffic-class-map dscp2tcFab
    dscp 5 traffic-class 2
    dscp 2 traffic-class 3
  !

emark-map tc2dscpFab
    traffic-class 0 dscp 63
    traffic-class 1 dscp 0
  !

remark-map tc2pcpFab
    traffic-class 2 cos 5
```

```
    traffic-class 1 cos 2
  !
```

- Interface Configuration

```
interface ethernet 0/3
  input
    trust dscp
    traffic-class-map dscp2tcFab
  !
  output
    remark dscp
    remark-map tc2dscpFab
  !
!
```

Fabric Interface (FabricInf) QoS Profile and Configuration (SLX OS vs OS ONE)

Use this topic to learn about the QoS configuration applied at the Fabric Interface (FabricInf) level. It includes QoS map creation, service policy configuration, profile creation using XCO, and the corresponding SLX OS and OS ONE running-config equivalents.

QoS Map and Profile Creation (XCO CLI)

- The following QoS maps define PCP–TC, DSCP–TC, TC–PCP, and TC–DSCP translations for fabric interfaces:

```
efa policy qos map create --name pcp2tcFabIF --type pcp-tc-map --rule "pcp[3],tc[1]"

efa policy qos map create --name dscp2tcFabIF --type dscp-tc-map --rule "dscp[3],tc[1]"

efa policy qos map create --name tc2pcpFabIF --type tc-pcp-map --rule "pcp[3],tc[1]"

efa policy qos map create --name tc2dscpFabIF --type tc-dscp-map --rule
"dscp[3],tc[2],dp[1]"
```

- A fabric-interface–level service policy is created:

```
efa policy qos service-policy-map create --name servicepolicy2 --rule "strict-
priority[1],dwrr[10;10;10;10;10;0;50],class[default]"
```

- The QoS profile associating these maps is created as follows:

```
efa policy qos profile create --name profileFabIF --trust dscp --pcp-
tc pcp2tcFabIF --dscp-tc dscp2tcFabIF --tc-pcp tc2pcpFabIF --service-policy
name[servicepolicy2],dir[out]
```

QoS Profile Details

Attribute	Value
Profile Name	profileFabIF
Trust Mode	dscp
PCP → TC Map	pcp2tcFabIF
DSCP → TC Map	dscp2tcFabIF
TC → PCP Map	tc2pcpFabIF
TC → DSCP Map	tc2dscpFabIF
Service Policy Map	servicepolicy2 (dir: out)

SLX OS Running Configuration

- QoS Maps

```
qos map cos-traffic-class pcp2tcFabIF
 map cos 3 to traffic-class 1 drop-precedence 0
!

qos map traffic-class-cos tc2pcpFabIF
 map traffic-class 1 drop-precedence 0 to cos 3
!

qos map traffic-class-dscp tc2dscpFabIF
 map traffic-class 2 drop-precedence 1 to dscp 3
!

qos map dscp-traffic-class dscp2tcFabIF
 map dscp 3 to traffic-class 1 drop-precedence 0
!
```

- Service Policy

```
policy-map servicepolicy2
  class default
    scheduler strict-priority 1 dwrr 10 10 10 10 10 0 50
!
```

- Interface Configuration

```
interface Port-channel 64
 speed 40000
 description MCTPeerInterface
 service-policy out servicepolicy2
 qos traffic-class-dscp tc2dscpFabIF
 qos dscp-traffic-class dscp2tcFabIF
 ip address 10.20.20.3/31
 no shutdown
!
```

OS ONE Running Configuration

- QoS Maps and Remark Maps

```
traffic-class-map pcp2tcFabIF
  cos 3 traffic-class 1
!
traffic-class-map dscp2tcFabIF
  dscp 3 traffic-class 1
!
remark-map tc2dscpFabIF
  traffic-class 2 dscp 3
!
remark-map tc2pcpFabIF
  traffic-class 1 cos 3
!
```

- Interface Configuration

```
interface port-channel 64
  input
    trust dscp
    traffic-class-map dscp2tcFabIF
  !
!
```

Tenant-Level QoS Profile and Configuration (SLX OS vs OS ONE)

Use this topic to learn about the QoS configuration applied at the Tenant level, including QoS map creation, service policy definition, profile configuration through XCO, and the corresponding outputs on SLX OS and OS ONE.

QoS Map and Profile Creation (XCO CLI)

- The following QoS maps define the PCP–TC, DSCP–TC, TC–PCP, and TC–DSCP mappings for the Tenant:

```
efa policy qos map create --name pcp2tcTen --type pcp-tc-map --rule "pcp[1],tc[2]"

efa policy qos map create --name dscp2tcTen --type dscp-tc-map --rule "dscp[1],tc[2]"

efa policy qos map create --name tc2pcpTen --type tc-pcp-map --rule "pcp[1],tc[2]"

efa policy qos map create --name tc2dscpTen --type tc-dscp-map --rule
"dscp[63],tc[0],dp[1]"
```

- A Tenant-level service policy is configured as follows:

```
efa policy qos service-policy-map create --name servicepolicy3 --rule "strict-
priority[2],dwrr[10;10;80;0;0;0],class[default]"
```

- The QoS profile referencing the above maps is created with:

```
efa policy qos profile create --name profileTen --trust dscp --pcp-tc pcp2tcTen
--dscp-tc dscp2tcTen --tc-dscp tc2dscpTen --tc-pcp tc2pcpTen --service-policy
name[servicepolicy3],dir[out]
```

QoS Profile Details

Attribute	Value
Profile Name	profileTen
Trust Mode	dscp
PCP → TC Map	pcp2tcTen
DSCP → TC Map	dscp2tcTen
TC → PCP Map	tc2pcpTen
TC → DSCP Map	tc2dscpTen
Service Policy Map	servicepolicy3 (dir: out)

SLX OS Running Configuration

- QoS Maps

```
qos map cos-traffic-class pcp2tcTen
 map cos 1 to traffic-class 2 drop-precedence 0
!

qos map traffic-class-cos tc2pcpTen
 map traffic-class 1 drop-precedence 0 to cos 3
!

qos map traffic-class-dscp tc2dscpTen
 map traffic-class 0 drop-precedence 1 to dscp 63
!

qos map dscp-traffic-class dscp2tcTen
```

```
 map dscp 1 to traffic-class 2 drop-precedence 0
!
```

- Service Policy

```
policy-map servicepolicy2
  class default
    scheduler strict-priority 1 dwrr 10 10 10 10 10 0 50
!
```

Interface-Level Configuration (SLX OS)

- Port-channel 1

```
interface Port-channel 1
 description EFA Port-channel pob31
 cluster-client auto
 switchport
 switchport mode access
 switchport access vlan 1102
 service-policy out servicepolicy3
 qos trust dscp
 qos remark dscp
 qos traffic-class-dscp tc2dscpTen
 qos dscp-traffic-class dscp2tcTen
 no shutdown
!
```

- Port-channel 2

```
interface Port-channel 2
 description EFA Port-channel pob32
 cluster-client auto
 switchport
 switchport mode trunk
 switchport trunk allowed vlan add 1401-1402
 no switchport trunk tag native-vlan
 service-policy out servicepolicy3
 qos trust dscp
 qos remark dscp
 qos cos-traffic-class pcp2tcTen
 qos traffic-class-dscp tc2dscpTen
 qos traffic-class-cos tc2pcpTen
 qos dscp-traffic-class dscp2tcTen
 no shutdown
!
```

- Port-channel 3

```
interface Port-channel 3
 description EFA Port-channel pob33
 cluster-client auto
 switchport
 switchport mode trunk-no-default-native
 switchport trunk allowed vlan add 1101
 service-policy out servicepolicy3
 qos trust dscp
 qos remark dscp
 qos cos-traffic-class pcp2tcTen
 qos traffic-class-dscp tc2dscpTen
 qos dscp-traffic-class dscp2tcTen
 no shutdown
!
```

- Port-channel 4 (partial)

```
interface Port-channel 4
 description EFA Port-channel pob34
 qos traffic-class-dscp tc2dscpTenIF
```

```
 qos dscp-traffic-class dscp2tcTenIF
 no shutdown
!
```

OS ONE Running Configuration

- QoS Maps and Remark Maps

```
traffic-class-map pcp2tcTen
  cos 1 traffic-class 2
!
traffic-class-map dscp2tcTen
  dscp 1 traffic-class 2
!
remark-map tc2dscpTen
  traffic-class 0 dscp 63
!
remark-map tc2pcpTen
  traffic-class 1 cos 3
!
```

Interface-Level Configuration (OS ONE)

- Ethernet 0/5

```
interface ethernet 0/5
  input
    trust dscp
    traffic-class-map dscp2tcTen
  !
  output
    remark dscp
    remark-map tc2dscpTen
  !
!
```

- Port-channel 2

```
interface port-channel 2
  input
    trust dscp
    traffic-class-map dscp2tcTen
  !
  output
    remark dscp
    remark-map tc2dscpTen
  !
!
```

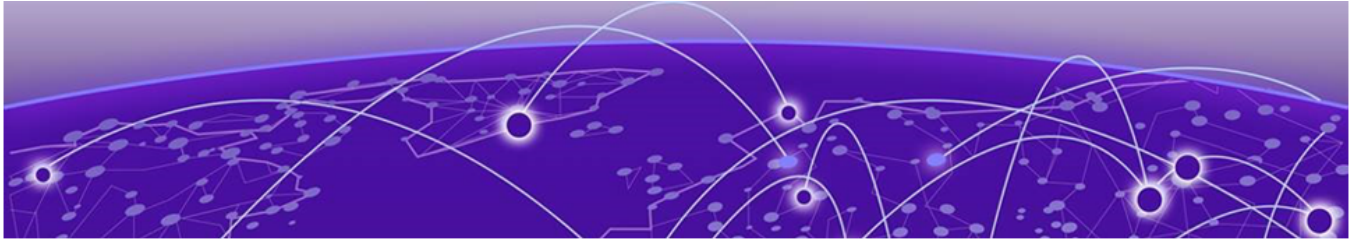

XCO Device Management

Learn about managing and configuring XCO and SLX devices.

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, 4.1.0*.
- 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.

- Extreme 8720--48XT
- Extreme 8520--48Y
- SLX 9540
- SLX 9640
- 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.

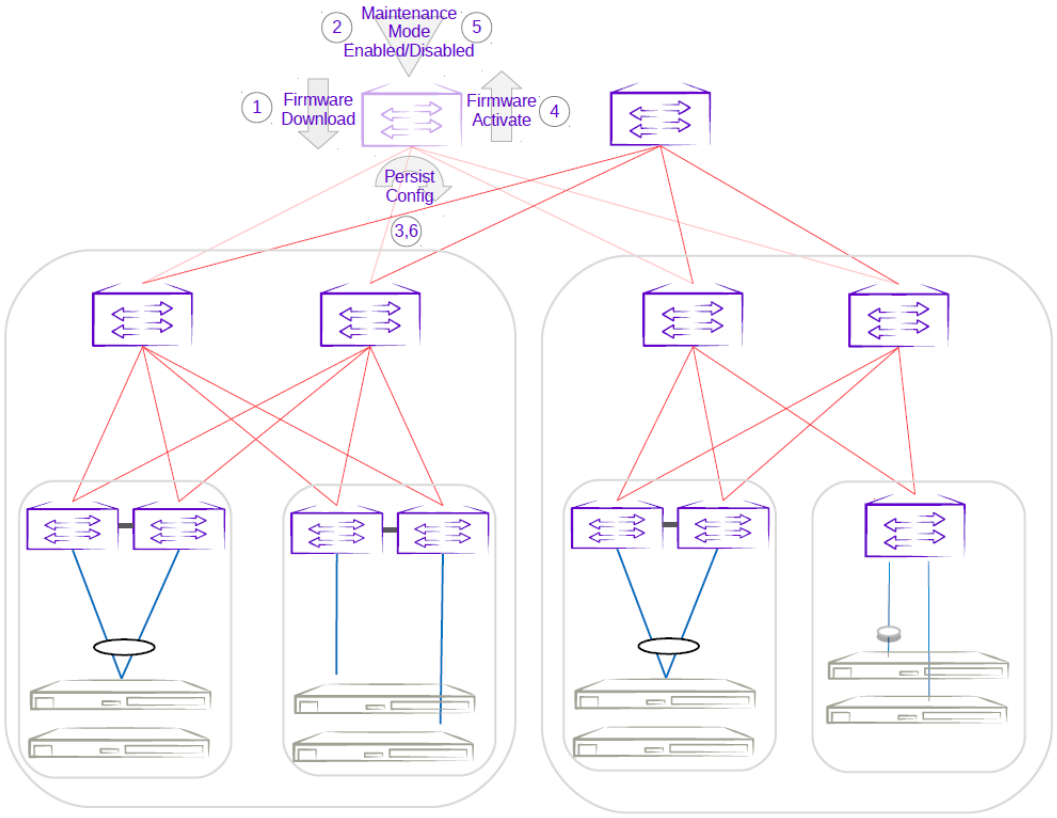

Figure 33: First super-spine firmware upgrade with maintenance mode

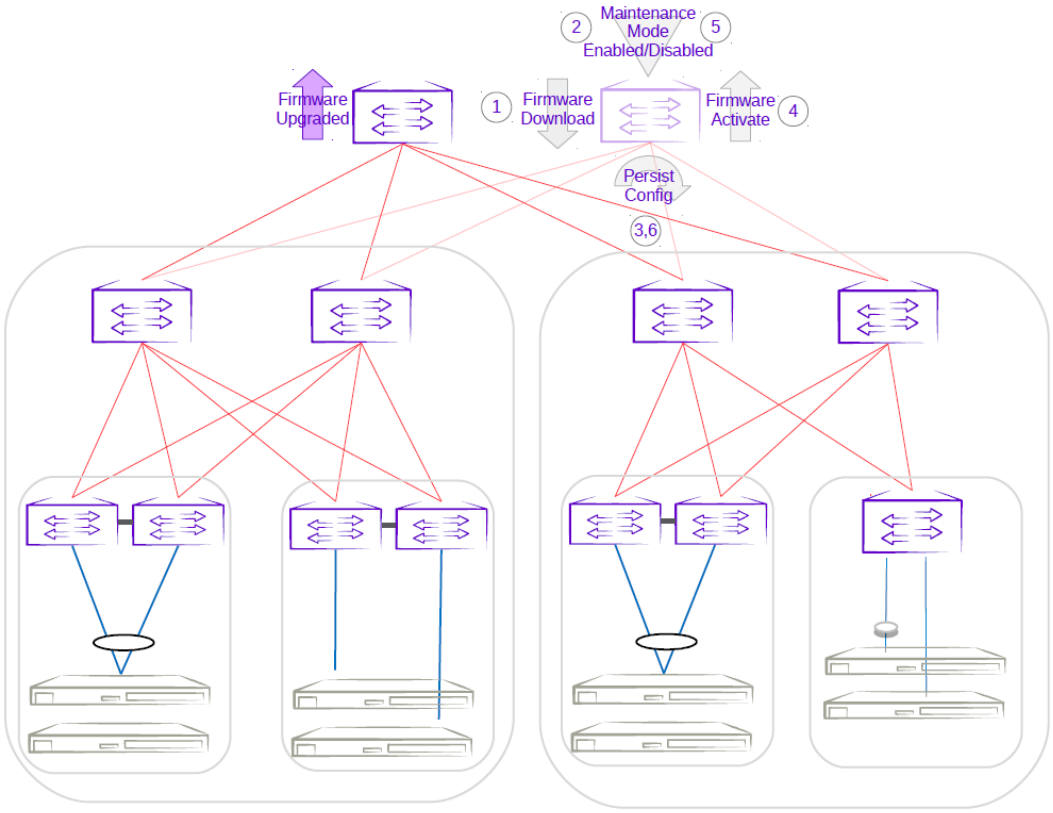

Figure 34: Second super-spine firmware upgrade with maintenance mode

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.

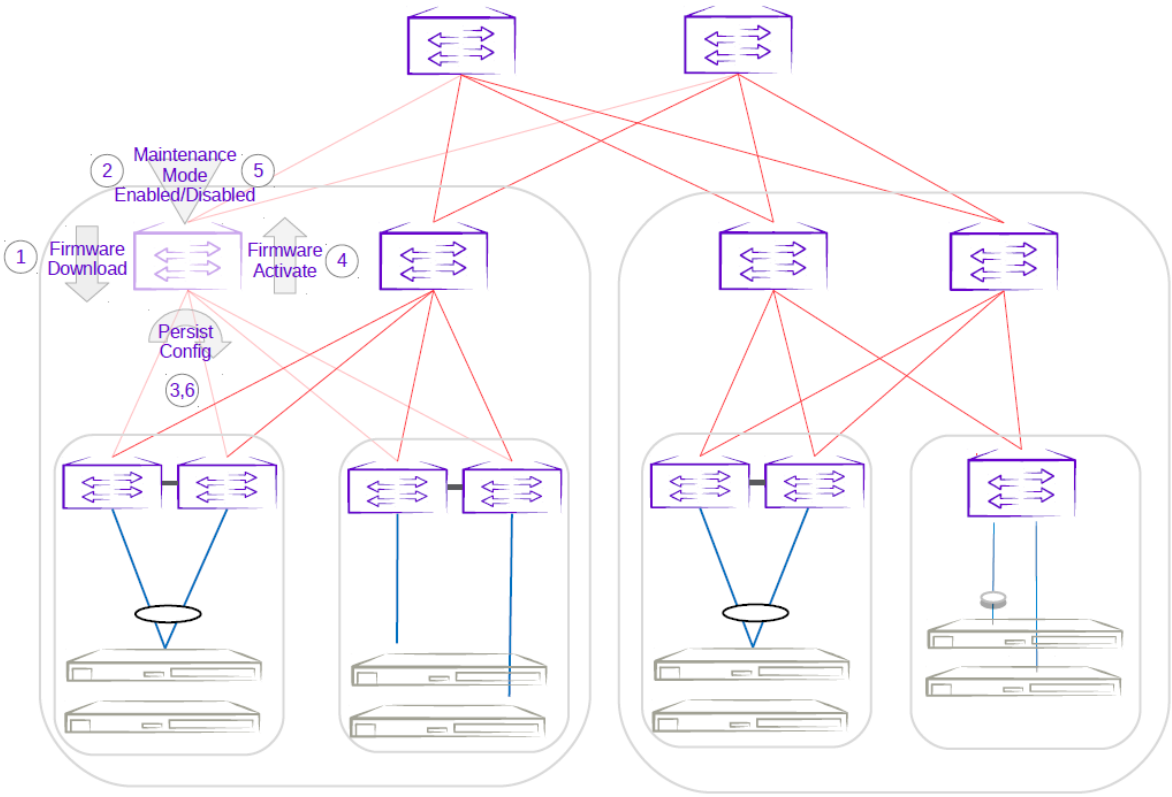

Figure 35: First spine firmware upgrade with maintenance mode

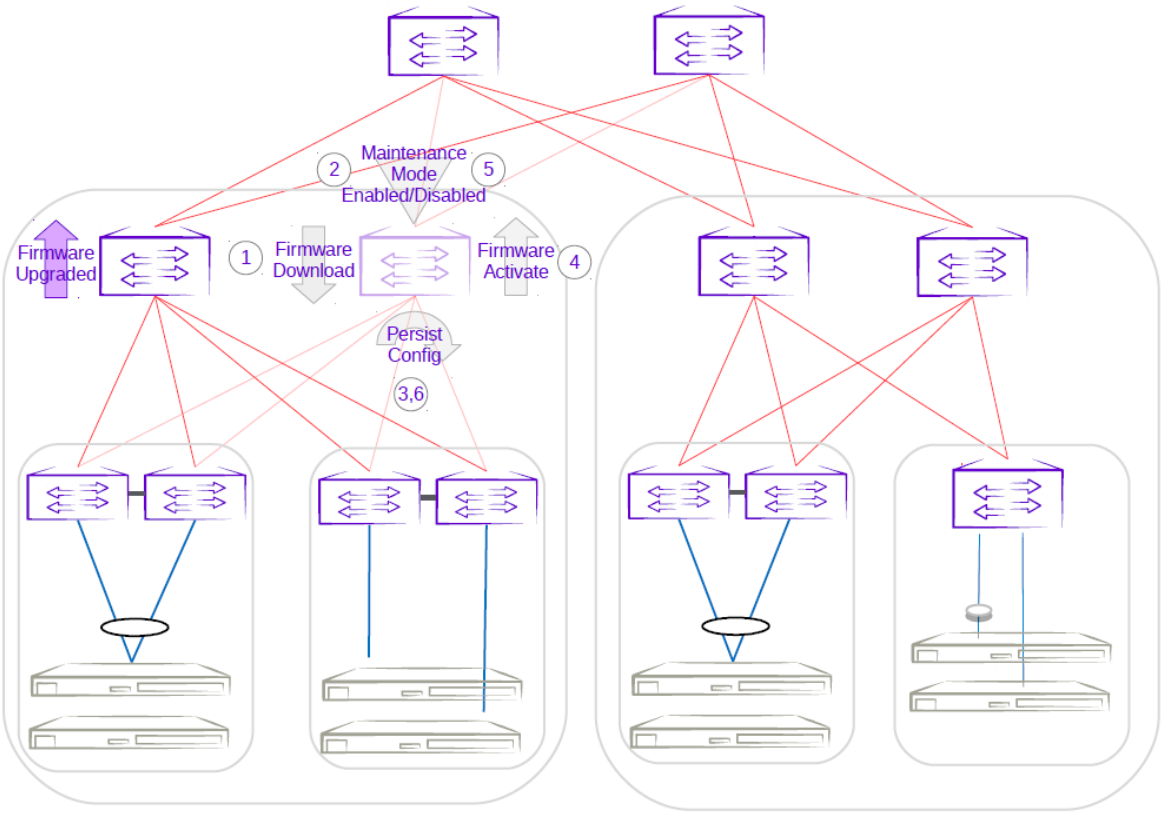

Figure 36: Second spine firmware upgrade with maintenance mode

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.

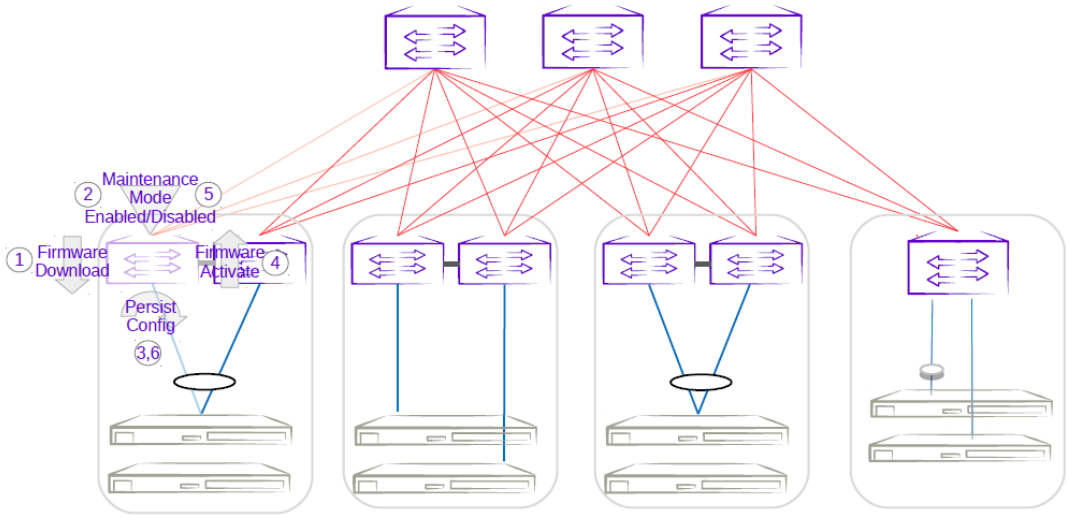

Figure 37: First MCT leaf firmware upgrade with maintenance mode

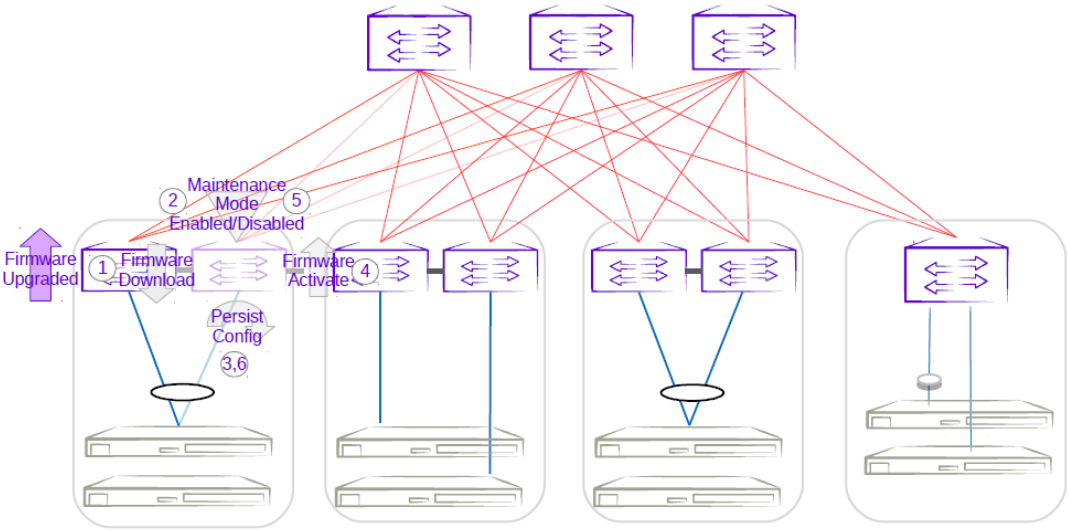

Figure 38: Second MCT leaf firmware upgrade with maintenance mode

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.

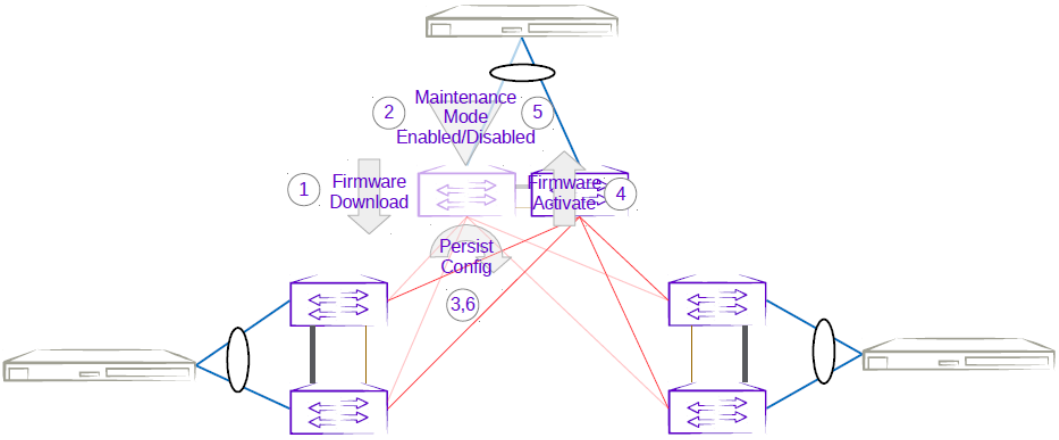

Figure 39: Three-rack centralized first MCT leaf firmware upgrade with maintenance mode

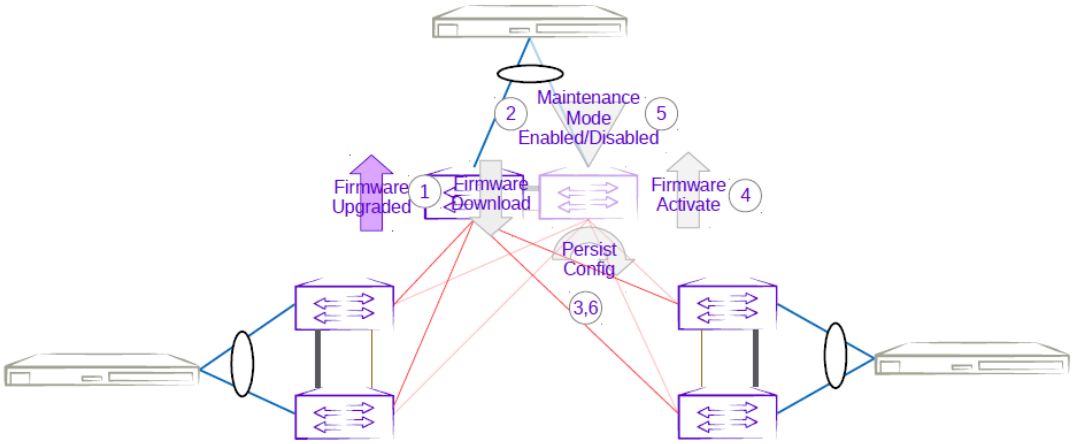

Figure 40: Three-rack centralized second MCT leaf firmware upgrade with maintenance mode

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.

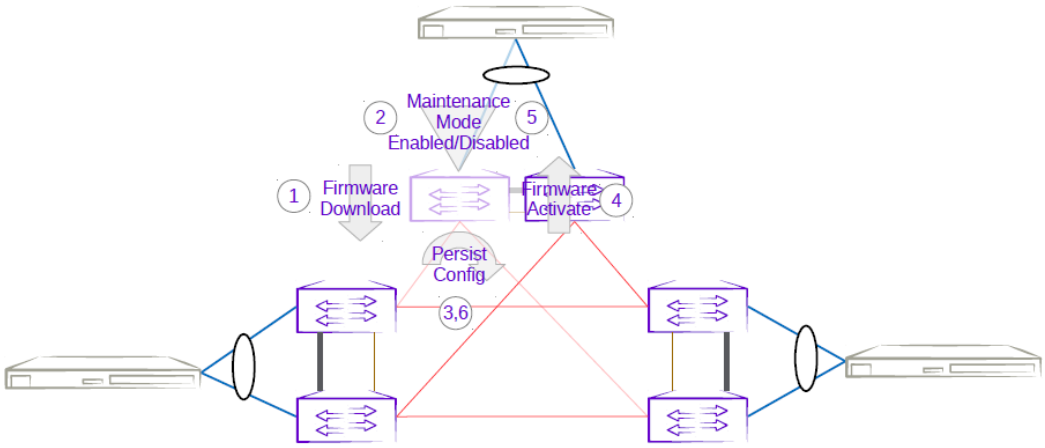

Figure 41: Three-rack ring first MCT leaf firmware upgrade with maintenance mode

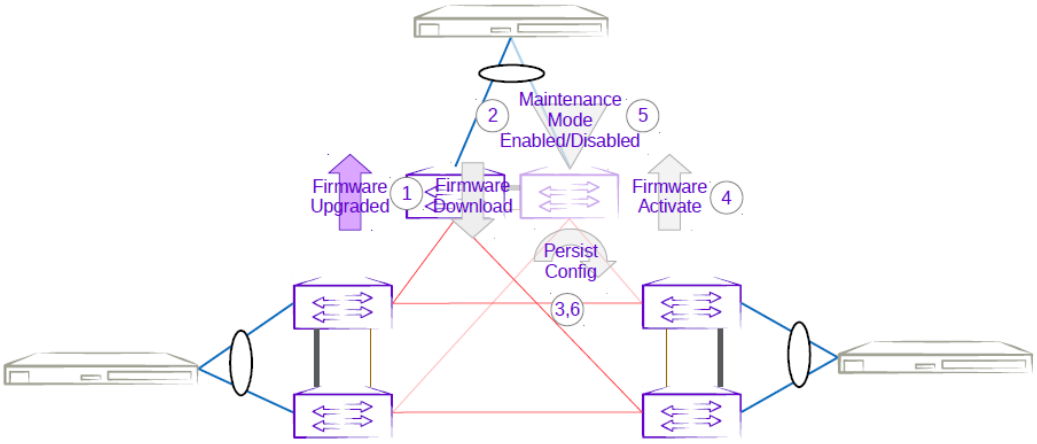

Figure 42: Three-rack ring second MCT leaf firmware upgrade with maintenance mode

Firmware Upgrade on OS ONE Devices

Use this topic to learn about the firmware upgrades on OS ONE devices through XCO, ensuring zero traffic loss across various fabric types (for example, non-Clos, multi-rack, 3-Clos, and 5-Clos).

The implementation supports firmware image upgrades for OS ONE devices using `.app` and `.incr` image formats, including full-install upgrade workflows initiated through XCO.

This topic covers the following points:

1. Upgrade Scope: Only the OS ONE to OS ONE firmware image upgrade is supported. Support for the SLX-to-SLX firmware upgrade is already available.
2. Upgrade Options: Only combined OS and application upgrades (OS+App) in *.bin format are supported.
3. Customizable Upgrade Flow: Allowing users to define their upgrade process . Flexible upgrade workflows with download, activate, and commit operations
4. Upgrade Monitoring: Tracking firmware upgrade status and maintaining historical data
5. Rollback: Automatic rollback on failed upgrades
6. Firmware changes: Historical tracking of all firmware changes
7. Fabric-wide upgrade: Support for fabric-wide upgrades with intelligent device grouping

Note
XCO firmware upgrade commands (CLI/REST) remain consistent with the XCO 3.7.0 release; refer to the existing command syntax and REST query structures.

Supported Deployment Topologies

The initial implementation focuses on the following deployment models:

- Single-rack deployments
- Non-Clos fabric environments

Future enhancements are expected to extend support to the following:

- Multi-rack deployments
- 3-stage Clos fabrics
- 5-stage Clos fabrics

Supported Image Types

OS ONE firmware updates use different image file types, each with distinct characteristics and requirements:

Image Type	File Extension	Reboot Required	Best Use Case
Full Operating System + Services	.bin	Yes	Major version upgrades or platform changes
Application Services Only	.app	No	Service patches without OS changes
Incremental Operating System	.incr	No	Base OS updates with minimal impact
Full Install (OS + Services)	full-install	Yes	Complete system replacement from backup

Firmware Transfer Support

The firmware image can be transferred to the device using the following protocols:

- SCP
- SFTP
- HTTP
- HTTPS

Note
XCO supports only the SCP protocol for firmware image transfer.

Best Practices

- **Schedule Off-Peak Windows –** Perform firmware upgrades during maintenance windows when network impact is acceptable
- **Test Non-Critical Devices First –** Upgrade a test or non-critical device before deploying to production
- **Use Application Updates for Services –** Deploy service patches with .app images when possible to avoid reboots
- **Verify Pre-Upgrade Status –** Confirm all devices are healthy before starting fabric-wide upgrades
- **Monitor During Upgrade –** Keep network monitoring active during upgrades to detect any issues promptly
- **Document Changes –** Record all firmware upgrades with dates, versions, and any observed impacts
- **Keep Firmware Organized –** Maintain a centralized repository of firmware images with clear naming conventions
- **Establish Rollback Criteria –** Define criteria for automatic rollback decisions in your environment

Firmware Upgrade Options Based on Image Types (OS ONE to OS ONE)

Use this topic to learn about the firmware upgrade comparison between SLX and OS ONE, OS ONE image types, and the sequence of CLI commands, UI flows, and options based on image extension.

CLI Commands for Firmware Upgrade on OS ONE Devices

Task	Command
Register firmware server	efa inventory firmware-host register --ip <ip> --protocol scp
Prepare .bin upgrade	efa inventory device firmware-download prepare add --fabric <fabric>
Prepare .app upgrade	efa inventory device firmware-download prepare add --ip <ips>
Execute upgrade	efa inventory device firmware-download execute --fabric <fabric>

Task	Command
Activate firmware	efa inventory device firmware-download activate --fabric <fabric>
Commit upgrade	efa inventory device firmware-download commit --fabric <fabric>
Rollback firmware	efa inventory device firmware-download rollback --fabric <fabric>
View upgrade status	efa inventory device firmware-download show --fabric <fabric>
View upgrade history	efa inventory device firmware-download history --fabric <fabric>

Customized Upgrade Flows

XCO supports multiple upgrade execution models to provide operational flexibility.

Note
As a best practice, complete the end-to-end flow cycle before triggering the next set of flows.

Option	Steps
Flow 1 — Sequential Upgrade Workflow	1. Download firmware image 2. Activate firmware 3. Commit or rollback firmware This workflow provides granular administrative control over each stage of the upgrade process.
Flow 2 — Combined Download and Activation	1. Download and activate firmware 2. Commit or rollback firmware This workflow reduces operational steps while preserving rollback flexibility.

Option	Steps
Flow 3 — Automated End-to-End Upgrade	1. Download firmware 2. Activate firmware 3. Commit firmware This workflow reduces operational steps while preserving rollback flexibility.
Flow 4 — Deployment-Based Upgrade Execution	Inventory-Based Upgrade • Select specific devices from inventory • Upgrade execution occurs in parallel • Traffic disruption may occur during the upgrade Fabric-Wide Upgrade • Select the target fabric • Devices are upgraded group-wise based on image type • Upgrade sequencing minimizes traffic disruption **Important:** Complete the entire upgrade cycle before initiating additional upgrade workflows.

Upgrade Firmware on OS ONE Devices

You can upgrade firmware on OS ONE devices.

Before You Begin

Before upgrading device firmware, ensure the following:

- Device is registered in XCO and assigned to a fabric
- Firmware image file is available on a local or remote server
- Firmware host (SCP/SFTP server) is accessible from the device
- Sufficient disk space is available on the device for new firmware
- Device is in a stable operational state

About This Task

Follow this procedure to upgrade firmware on OS ONE devices.

Procedure

1. Register firmware host.

   Configure XCO with the location of your firmware server. XCO will retrieve firmware images from this location during upgrades:

   ```
   efa inventory firmware-host register --ip <server-ip> --protocol scp --username
   <username> --password <password>
   ```

   Example: Register firmware server at 10.32.89.26 for uploading firmware images via SCP

2. Upgrade full OS and Services (.bin).

   Use .bin images for major upgrades involving both operating system and service components. These upgrades require device reboot and support fabric-wide deployment with intelligent grouping to avoid traffic loss.

   a. Prepare firmware for fabric-wide upgrade.

   ```
   efa inventory device firmware-download prepare add --fabric <fabric-name> --
   firmware-directory <path>/<filename>.bin --firmware-host <server-ip>
   ```

   b. Run the upgrade with automatic maintenance mode and DRC enabled.

   ```
   efa inventory device firmware-download execute --fabric <fabric-name> --noActivate
   --noAutoCommit --drc
   ```

   c. Activate the firmware.

   ```
   efa inventory device firmware-download execute --fabric <fabric-name> --noActivate
   --noAutoCommit --drc
   ```

   d. Commit the upgraded firmware.

   ```
   efa inventory device firmware-download commit --fabric <fabric-name>
   ```

3. Upgrade services only (.app).

   Use .app images for service-level updates that do not require system reboot. Application upgrades can be deployed in parallel across all devices with no traffic loss.

   a. Prepare firmware update:

   ```
   efa inventory device firmware-download prepare add --ip <device-ip-list> --firmware-
   directory <path>/<filename>.app --firmware-host <server-ip> --prepared-list-name
   <list-name>
   ```

   b. Run parallel upgrade.

   ```
   efa inventory device firmware-download execute --prepared-list-name <list-name> --
   noActivate --noAutoCommit --noMaintMode
   ```

4. Upgrade OS incrementally (.incr).

   Use .incr images for incremental base OS updates that do not require reboot. Like application upgrades, incremental updates can run in parallel.
   Prepare incremental upgrade:

   ```
   efa inventory device firmware-download prepare add --ip <device-ip-list> --firmware-
   directory <path>/<filename>.incr --firmware-host <server-ip> --prepared-list-name
   <list-name>
   ```

5. Full install upgrade.

   Full install upgrades replace the entire system with a new firmware build from your server. This process requires explicit marking and is useful for emergency recoveries.

   a. Prepare full install upgrade.

   ```
   efa inventory device firmware-download prepare add --ip <device-ip> --firmware-
   directory <path>/<filename>.bin --firmware-host <server-ip> --full-install
   ```

   b. Run full install.

   ```
   efa inventory device firmware-download execute --fabric <fabric-name> --noActivate
   --drc
   ```

Monitor Firmware Upgrades

You can monitor the firmware upgrade status and maintain the historical data. The commands show the OS ONE and switching image version.

About This Task
Follow this procedure to monitor the firmware upgrade status and maintain the historical data.

Procedure

1. View upgrade status.
   Display the status of current and previous firmware upgrades:
   ```
   efa inventory device firmware-download show --fabric <fabric-name>
   ```
2. View upgrade history.
   Review the history of firmware upgrades on devices:
   ```
   efa inventory device firmware-download history --fabric <fabric-name>
   ```
3. View operational history.
   Track firmware upgrade operations and associated events:
   ```
   efa inventory device firmware-download operational-history --fabric <fabric-name>
   ```
4. Device information.
   View current firmware versions on specific devices:
   ```
   efa inventory device list --ip <device-ip>
   ```

Maintenance Mode and DRC

Learn about maintenance mode and DRC.

Automatic Traffic Protection

For full OS upgrades requiring reboot, XCO automatically enables Maintenance Mode (MM) and Drift Reconciliation (DRC) to manage traffic and configuration consistency.

Maintenance Mode

Maintenance Mode suspends traffic on individual devices during reboot when they are members of a fabric. Fabric-wide upgrades are automatically grouped so that not all devices reboot simultaneously, preserving network connectivity and avoiding traffic loss.

Drift Reconciliation

After device reboot completes, DRC automatically verifies and reconciles any configuration differences between the device and the intended fabric configuration, ensuring consistency across the network.

Firmware Status Updates

The firmware status is updated based on the RASLog events.

```
"Firmware","LogID":5021,"Msg":"App SR Installed successful."}

"Firmware","LogID":8005,"Msg":"Failed to activate firmware SR-22.1.6.0-009.incr, err
Activation failed -  SR firmware should be committed before upgrade"}

"Firmware","LogID":8007,"Msg":"Commit current Base firmware image"}
```

```
"Firmware","LogID":5021,"Msg":"Firmware change successful.  Current Firmware Version is
ONEOS-10.1.3.0-009"}

Looking raslog event :"Msg":"After SR App Installation, all services in ready state"}
```

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, 4.1.0*

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 |Extreme 8720|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
<username> --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.

Firmware Download (FWDL) & Maintenance Mode (MM) Behavior in Extreme OS ONE

Learn about how Firmware Download (FWDL) in OS ONE interacts with Maintenance Mode (MM) settings on network devices, specifically, how MM and MM-ON-REBOOT flags are handled before, during, and after FWDL operations.

Note
FWDL must preserve the original Maintenance Mode state. What was enabled or disabled before FWDL must remain the same after FWDL completes.

Case 1: Switch Config > MM-Disabled, MM-ON-REBOOT-Disabled

- Option A: FWDL <device> --noMaintMode - MM configs untouched.
  - No temporary MM enable or disable cycle.
  - Final State: `MM = Disabled`, `MM-ON-REBOOT = Disabled`

- Option B: FWDL <device> (default behavior) - Temporary MM enablement for FWDL:
  1. MM-ON-REBOOT enabled
  2. MM enabled
  3. Firmware downloaded & activated
  4. MM disabled
  5. MM-ON-REBOOT disabled
- Final State: `MM = Disabled`, `MM-ON-REBOOT = Disabled`

Case 2: Switch Config > MM-Enabled, MM-ON-REBOOT (Enabled or Disabled) FWDL is BLOCKED if the device is already in Maintenance Mode (MM = Enabled)

Prevents FWDL from interfering with ongoing maintenance operations.

Case 3: Switch Config > MM-Disabled, MM-ON-REBOOT-Enable

- Option A: FWDL <device> --noMaintMode`- Device reboots to activate firmware → enters MM automatically.
  - FWDL flow explicitly disables MM post-reboot.
  - Final State: `MM = Disabled`, `MM-ON-REBOOT = Enabled`
- Option B: FWDL <device>` _(default behavior)_- MM → Enabled (temporarily)
  - Firmware downloaded & activated
  - MM disabled
- Final State: `MM = Disabled`, `MM-ON-REBOOT = Enabled`

FWDL Reboot Behavior Based on enable-on-reboot & FWDL Flags

enable-on-reboot (device flag)	noMaint (FWDL flag)	DRC (FWDL flag)	FWDL Reboot Behavior	OS ONE Final MM Behavior
True / False	False	False	• MM enabled before reboot • NO DRC run • MM disabled after reboot	• MM-ON_REBOOT enabled • MM enabled
True / False	False	True	• MM enabled before reboot • DRC triggered • MM disabled after DRC completes	• MM-ON_REBOOT enabled • MM enabled

enable-on-reboot (device flag)	noMaint (FWDL flag)	DRC (FWDL flag)	FWDL Reboot Behavior	OS ONE Final MM Behavior
True	True	False	• Device (switch) enters MM via enable_on_reboot during reboot • NO DRC • EFA disables MM post-reboot	• MM enabled during reboot • FWDL
False	True	False	• No MM activation • NO DRC	No MM change, stays OFF
True	True	True	• Device (switch) enters MM via enable_on_reboot during reboot • DRC triggered • MM disabled post-DRC	• MM enabled during reboot • FWDL
False	True	True	• Invalid use case. Errors are raised	NA CLI or API rejects this combo

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, 4.1.0*.

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.

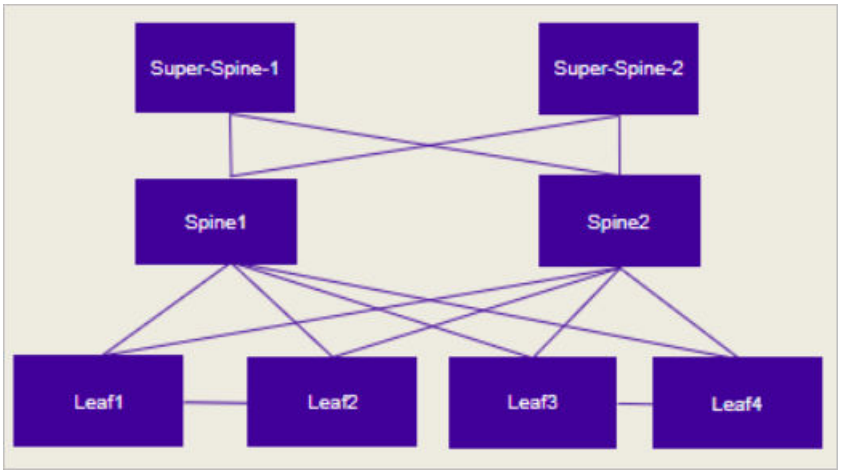

Figure 43: Batches for fabric-wide update

   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 --help
Execute firmware download for executed devices

Usage:
  efa inventory device firmware-download execute [flags]

Flags:
      --fabric string               Fabric
      --prepared-list-name string   Prepared list name
      --noAutoCommit                Configure Auto commit in Firmware Download
      --noMaintMode                 Configure Maintenance Mode in Firmware Download
      --noActivate                  Configure Activation in Firmware Download
      --drc                         Configure a drift reconciliation operation is
performed after Firmware Download
      --group-execution string      Configure Group Execution Policy <continue-on-
error | stop-on-error> in Firmware Download.
  --- Time Elapsed: 814.943µs ---
(efa:user)user@server2:~$
```

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.

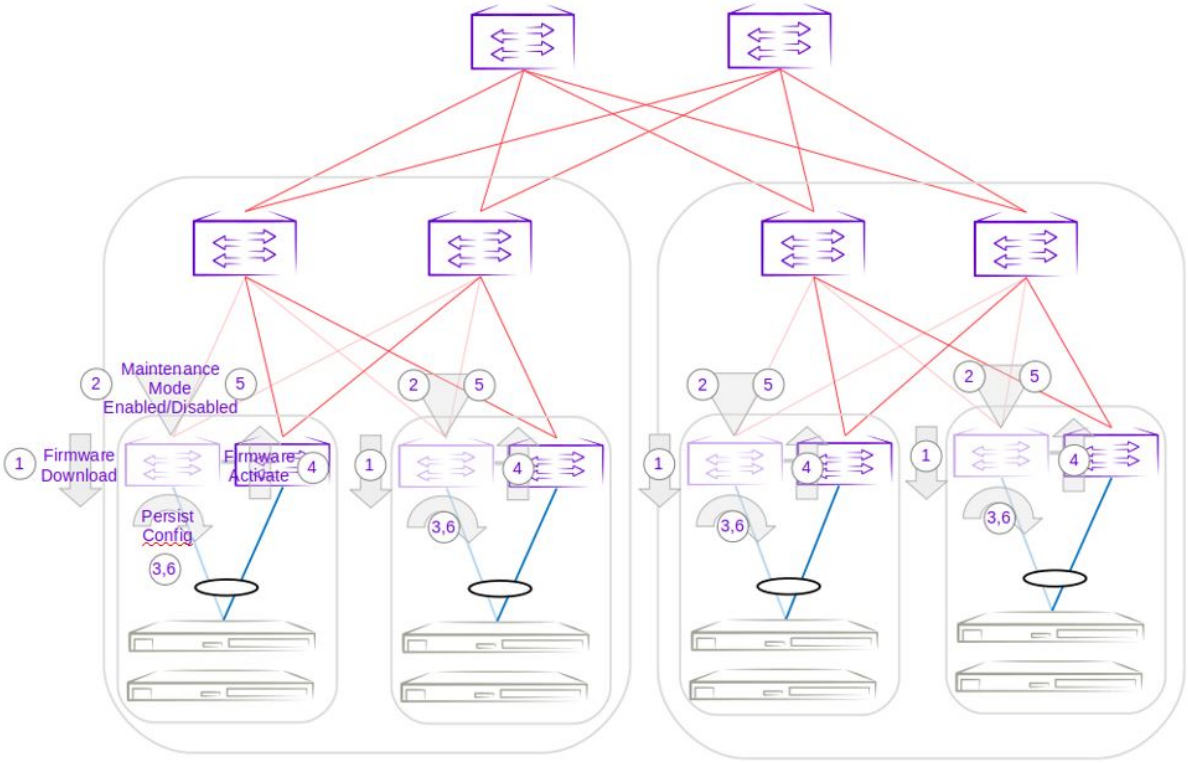

Figure 44: Group 1 - Firmware download execution of lower IP address MCT Peer leaf devices

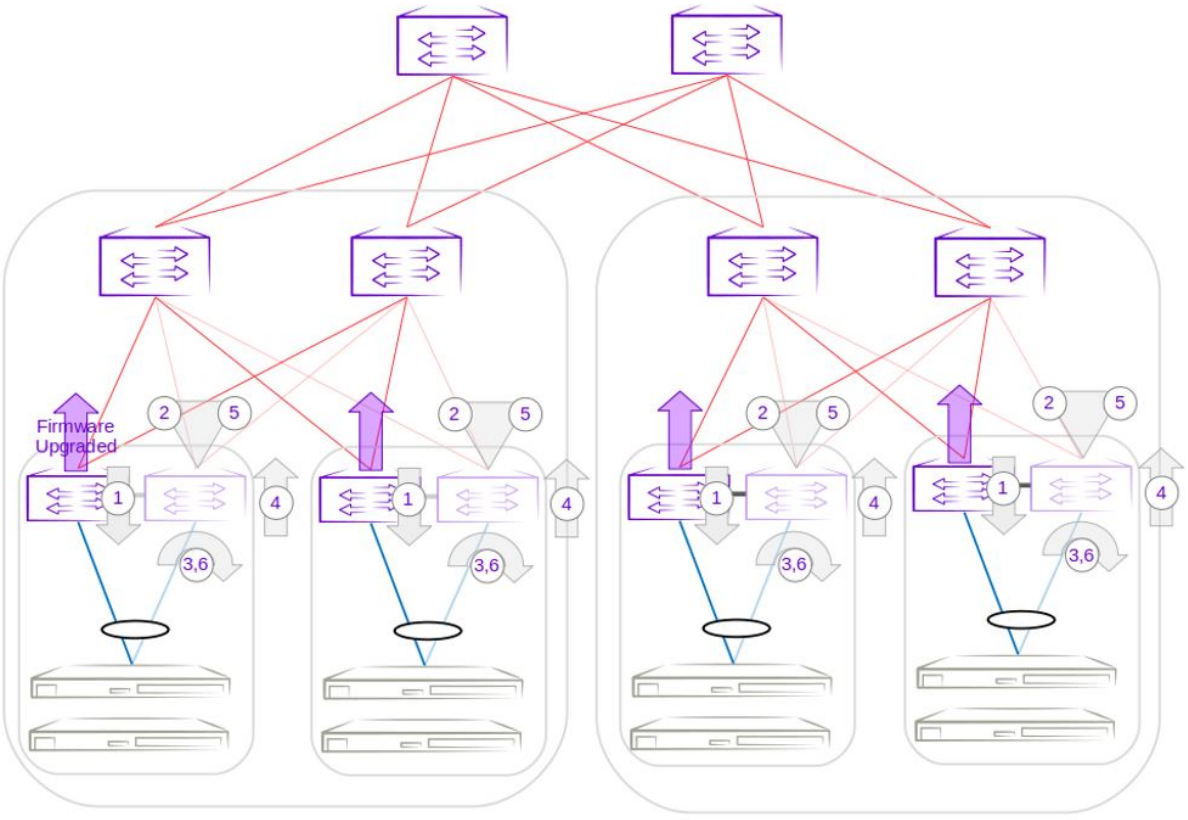

Figure 45: Group 2 - Firmware download execution of higher IP address MCT Peer leaf devices

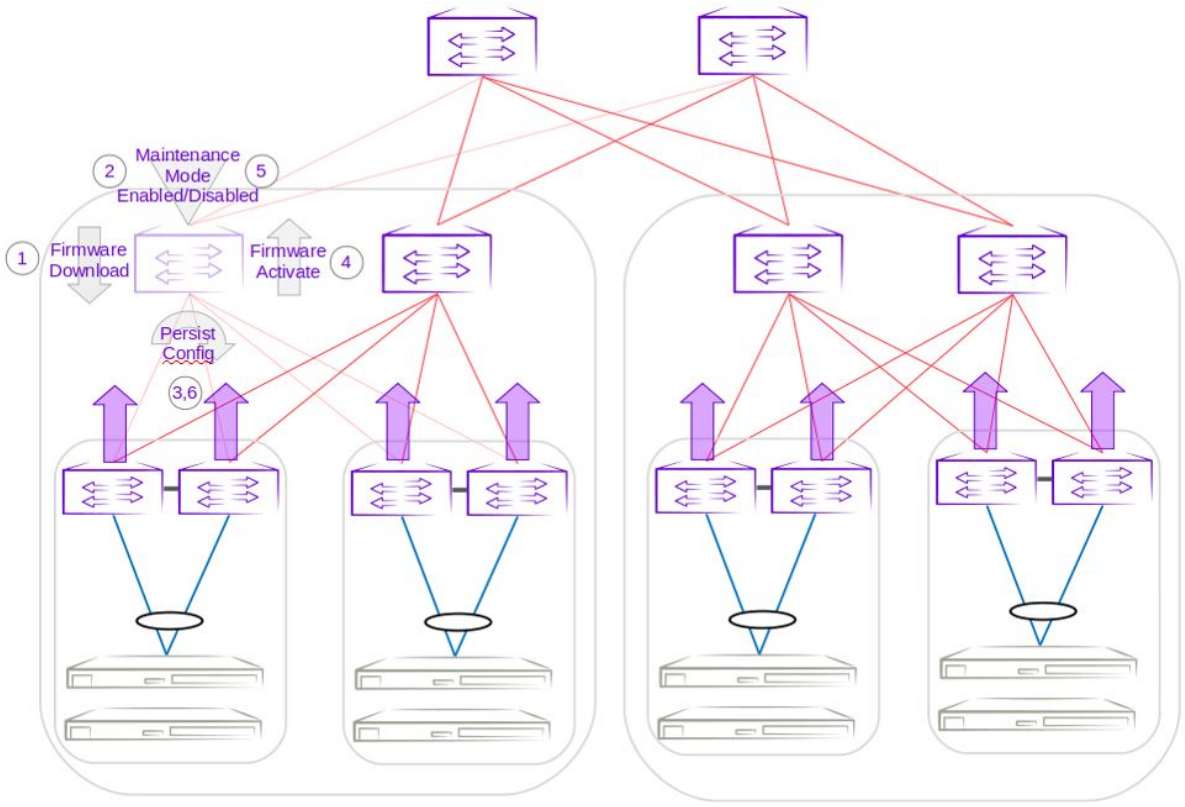

Figure 46: Group 3 - Firmware download execution of lowest IP address spine device

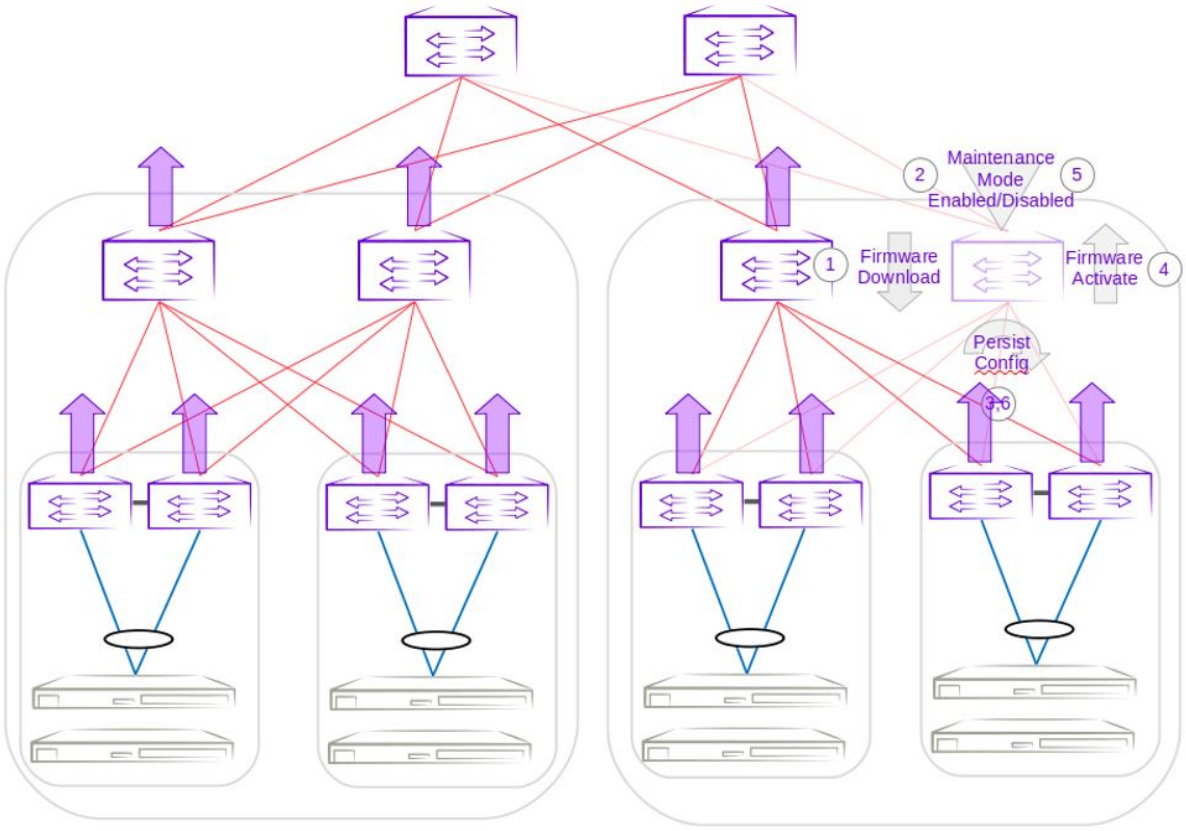

Figure 47: Group 4 - Firmware download execution of highest IP address spine device

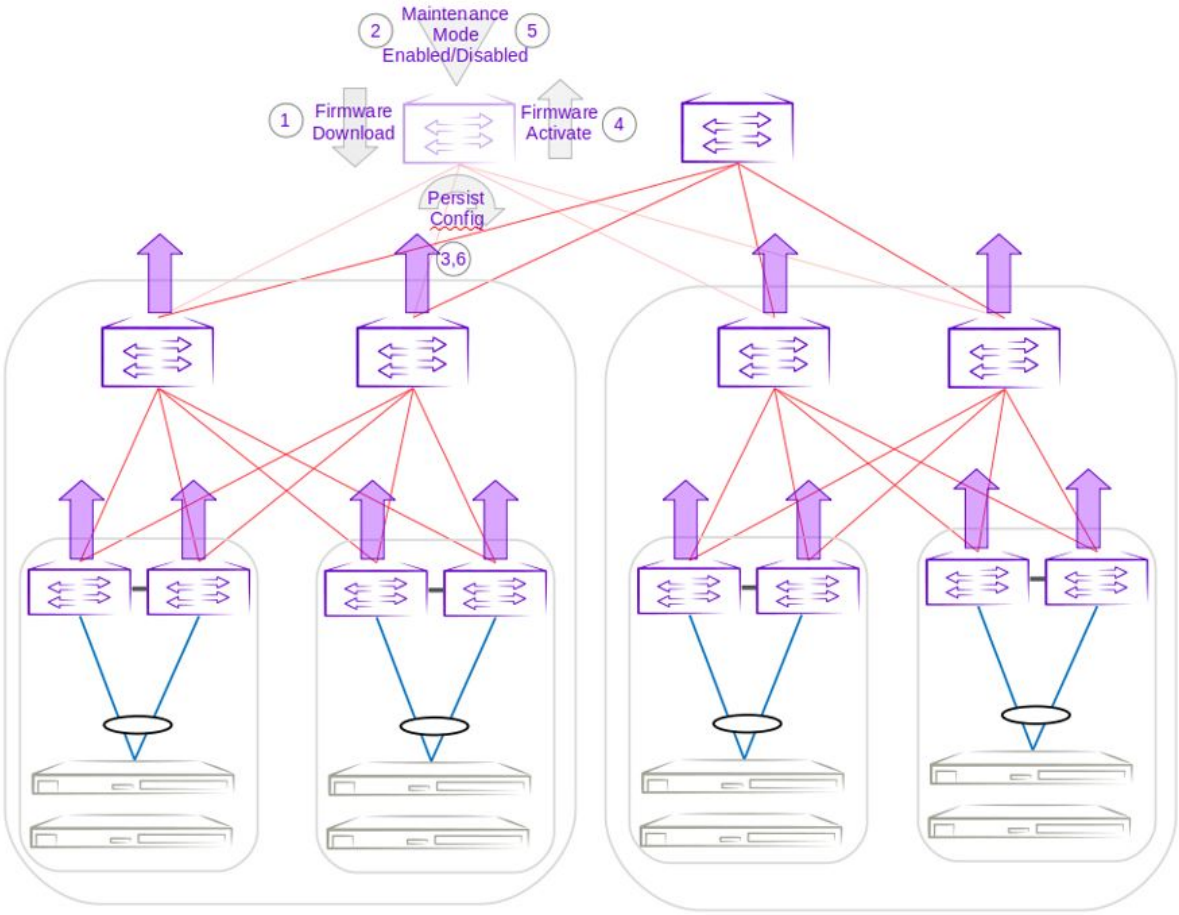

Figure 48: Group 5 - Firmware download execution of lowest IP address super-spine device

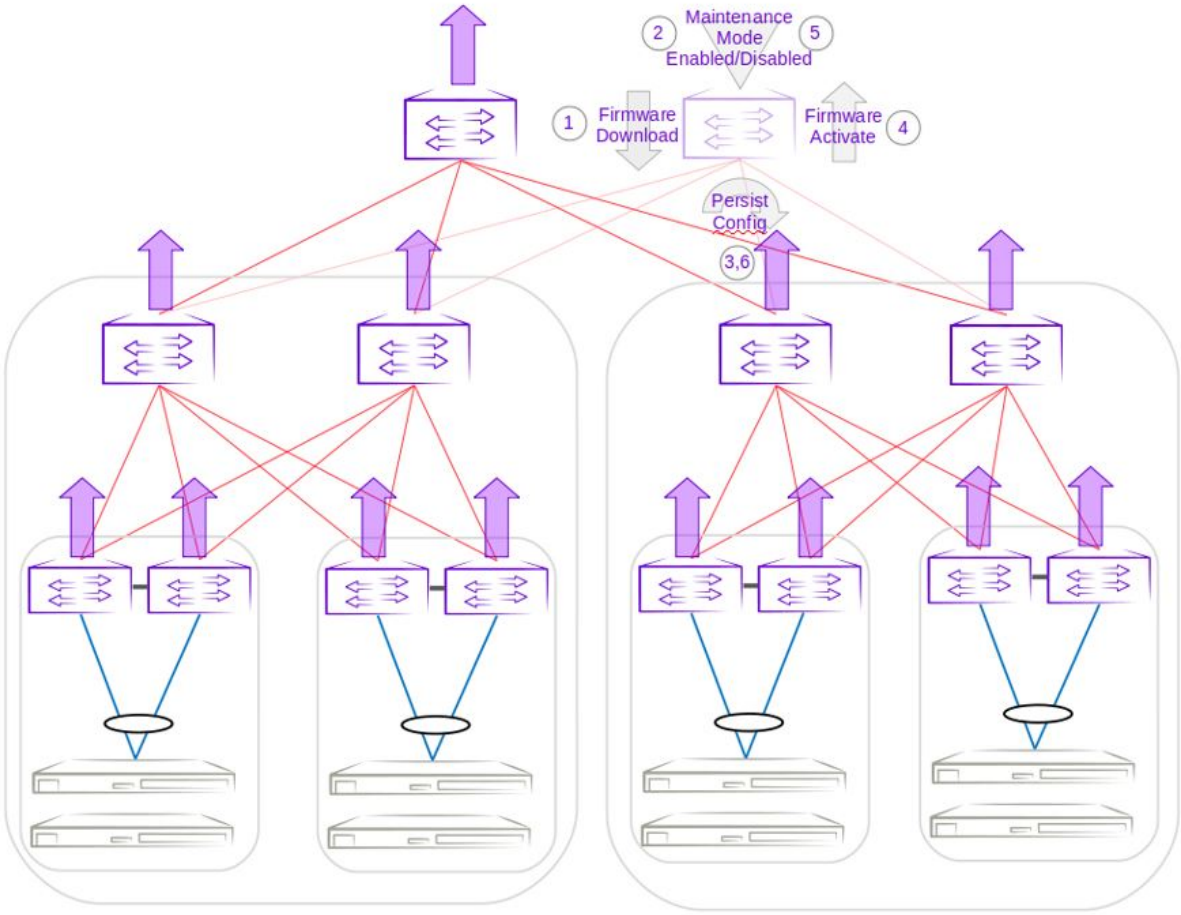

Figure 49: Group 6 - Firmware download execution of highest IP address super-spine device

Firmware Upgrade Status

You can fetch the firmware upgrade status using the CLI.

You can get the detailed overview of the upgrade process, including the start time, download completion time, and commit time. You can access the upgrade history either based on user action or runtime status.

Note
When you trigger a firmware download on an SLX device outside the XCO framework, the historical data might not be automatically stored in the database.

Use the `efa inventory device firmware-download` command to fetch the firmware upgrade history.

Get User Action Fabric-Wide Firmware-Download (Fwdl) History

You can fetch the firmware download history based on the execution order.

About This Task

Follow this procedure to get the list of devices for firmware download history details in execution order.

When you trigger a firmware upgrade action (Download, Activate, and Commit), the sequence of actions performed during the upgrade process is captured in the history output.

- The version transitioning from source version to target version
- Status updates during the upgrade process
- A list of execution IDs performed at the fabric level
- Details about how the group strategy was generated and which devices were part of the execution
- Execution Order based on show output entries

For information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

Run the following command to get the list of firmware download devices for a firmware-download history operation:

```
efa inventory device firmware-download history --fabric fabric1

Execution EntryNumber :7, UUID :d3874b29-3d33-4b57-a929-c4f0bfa40ab9
Execution Start :2024-03-14 17:03:42 -0700 PDT, End :2024-03-14 17:22:02 -0700 PDT,Duration
:18m19.570780811s
Previous Version :20.6.2slxos20.6.2_240313_0330, Target Version :20.5.3a
Force :true, FullInstall :false
Group Execution Policy :continue-on-error, Prepared-list-name :
+-------+-------------+-------------+-----------+------+-------------+--------------------
+--------------------------------+--------------------------------+-------------------------------
+------------------+------------+-----------------+
| Group |  IP Address  | User Task ID | Host Name | Role
| Update State |       Status       |          Start Time           |
End Time            |          Update Time          | User Input Flags | FabricName | Detailed Status |
+-------+-------------+-------------+-----------+------+-------------+--------------------
+--------------------------------+--------------------------------+-------------------------------
+------------------+------------+-----------------+
| 1     | 10.20.55.173 | 11           | B145-R173 | Leaf |
Completed    | Firmware Committed | 2024-03-14 17:03:42 -0700 PDT | 2024-03-14
17:22:02 -0700 PDT | 2024-03-14 17:21:29 -0700 PDT | D,A,C,MM         | fabric1    |                 |
+-------+-------------+-------------+-----------+------+-------------+--------------------
+--------------------------------+--------------------------------+-------------------------------
+------------------+------------+-----------------+
```

Get Runtime (Operational) Status-based History

You can fetch the firmware download history based on the runtime status.

About This Task

Follow this procedure to get the list of firmware download devices for a firmware download history operation.

When you trigger a firmware upgrade action, the sequence of actions performed during the upgrade process is captured in history.

- Details about the source firmware version and the target version
- Status updates throughout the upgrade process
- Granular level operations, including (MM enable or disable, download, activate, commit, DRC, persist config saved). Only SLX switch operation tasks are captured in the operational table, and EFA level logic tasks such as validation device status, validate firmware, queue operational are not captured
- Start and end times for each granular operation are calculated by EFA based on configure time as start time and polling status end result as end time
- For each execution, details about how the group strategy was generated and which devices were part of the execution. This includes information on devices involved, actions taken, strategies, and granular level of internal tasks implemented at the device level
- Execution order of show output entry details: <Execution-id> <User Action> <Group-id> <Device-ip> <Granular level Data>

For information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

Run the following command to get the list of firmware download devices for a firmware download history operation:

```
efa inventory device firmware-download operational-history --fabric fabric1

Execution EntryNumber :7, UUID :d3874b29-3d33-4b57-a929-c4f0bfa40ab9
Execution Start :2024-03-14 17:03:42 -0700 PDT, End :2024-03-14 17:22:02 -0700 PDT,Duration
:18m19.570780811s
Previous Version :20.6.2slxos20.6.2_240313_0330, Target Version :20.5.3a
Force :true, FullInstall :false
Group Execution Policy :continue-on-error, Prepared-list-name :
+-------+--------------+---------+--------------+-----------+------------------------------
+-------------------------------+-------------------------------+
| Group |  IP Address  | Oper Id | User Task Id | Host
Name |        Detail Status         |          Starte Time          |           End Time            |
+-------+--------------+---------+--------------+-----------+------------------------------
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 53      | 11           | B145-R173
| Firmware Commit              | 2024-03-14 17:21:14 -0700 PDT | 2024-03-14 17:21:29 -0700 PDT |
+-------+--------------+---------+--------------+-----------+------------------------------
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 51      | 11           | B145-R173
| Persist Config After Reload  | 2024-03-14 17:21:00 -0700 PDT | 2024-03-14 17:21:14 -0700 PDT |
+-------+--------------+---------+--------------+-----------+------------------------------
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 49      | 11           | B145-
R173 | Maintenance Mode Disable     | 2024-03-14 17:18:54 -0700 PDT | 2024-03-14 17:21:00 -0700 PDT |
+-------+--------------+---------+--------------+-----------+------------------------------
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 47      | 11           | B145-
R173 | Firmware Activate            | 2024-03-14 17:07:43 -0700 PDT | 2024-03-14 17:18:54 -0700 PDT |
+-------+--------------+---------+--------------+-----------+------------------------------
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 45      | 11           | B145-
R173 | Persist Config Before Reload | 2024-03-14 17:07:28 -0700 PDT | 2024-03-14 17:07:43 -0700 PDT |
+-------+--------------+---------+--------------+-----------+------------------------------
```

```
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 43      | 11           | B145-
R173 | Firmware Download            | 2024-03-14 17:04:00 -0700 PDT | 2024-03-14 17:07:28 -0700 PDT |
+-------+--------------+---------+--------------+-----------+-------------------------------
+-------------------------------+-------------------------------+
| 1     | 10.20.55.173 | 41      | 11           | B145-
R173 | Maintenance Mode Enable      | 2024-03-14 17:04:00 -0700 PDT | 2024-03-14 17:05:03 -0700 PDT |
+-------+--------------+---------+--------------+-----------+-------------------------------
+-------------------------------+-------------------------------+
```

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.

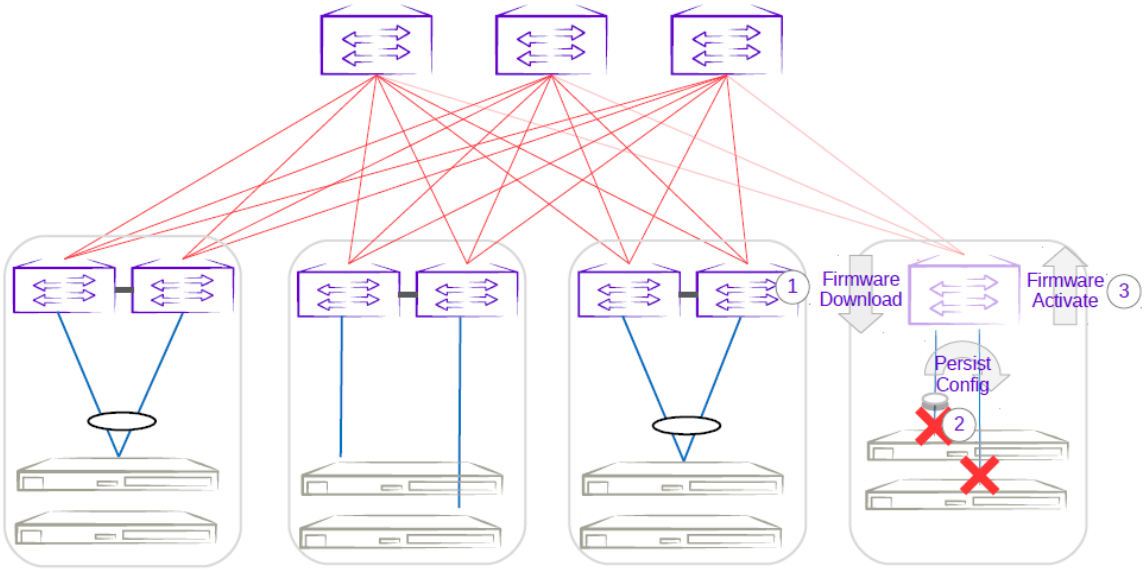

Figure 50: Single-leaf traffic loss

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.

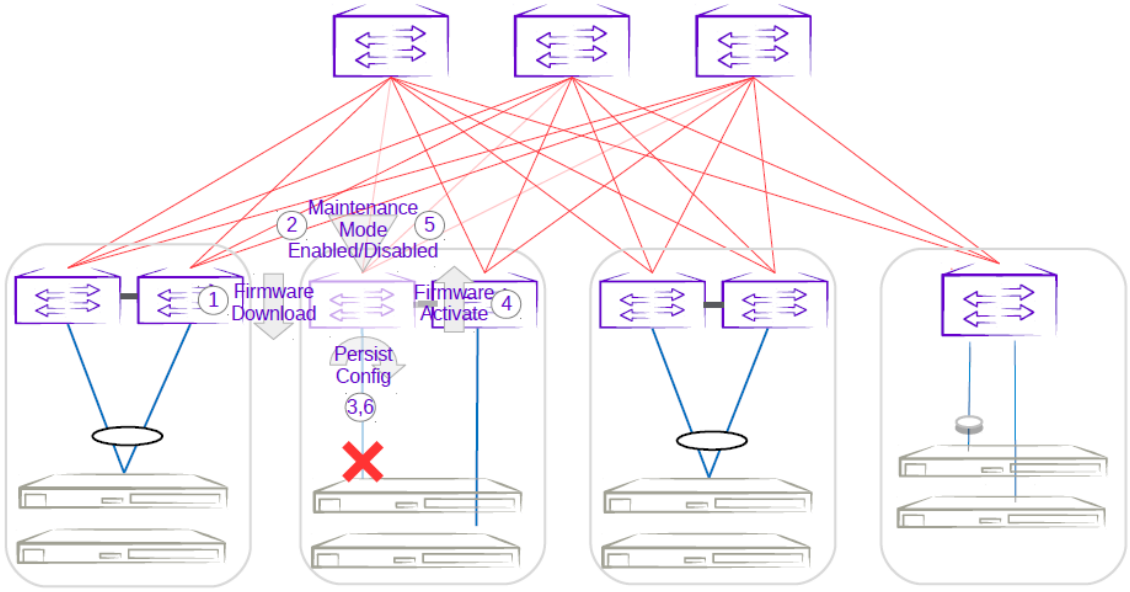

Figure 51: Single-homed server traffic loss

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.

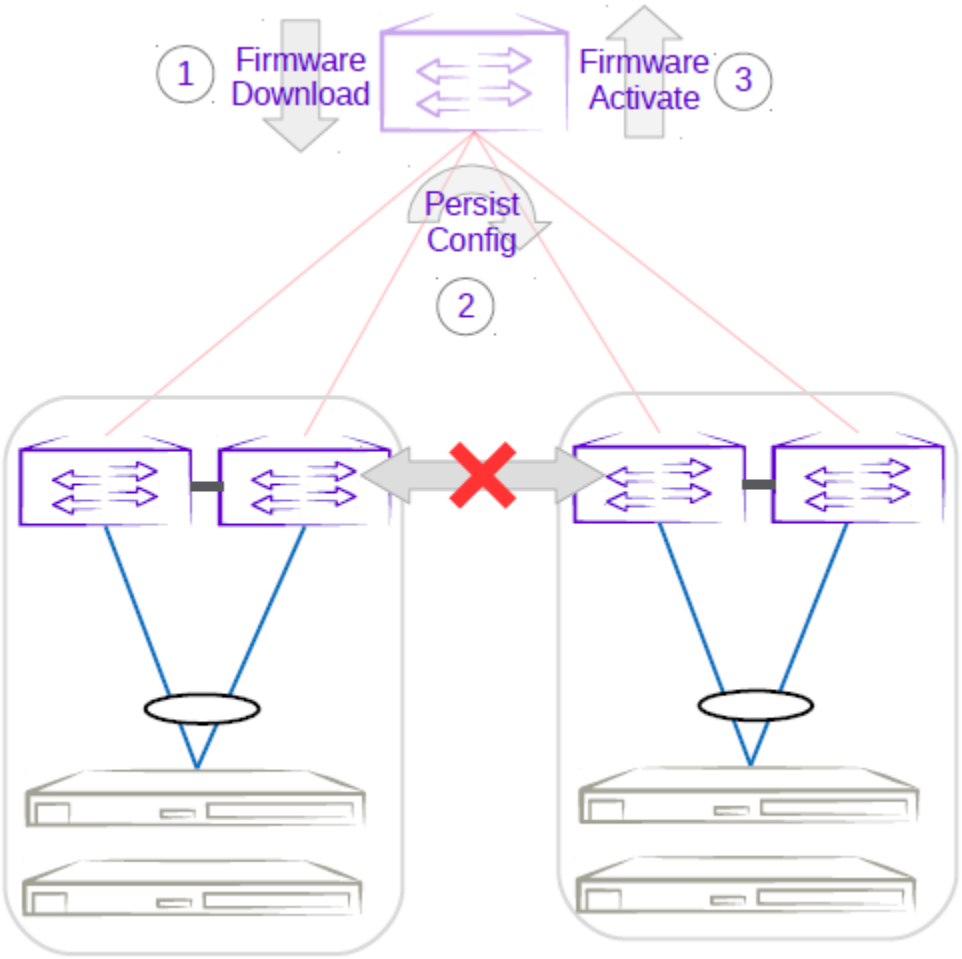

Figure 52: Non-redundant spine traffic loss

Firmware Upgrade Auto Recovery

Extreme Cloud Orchestration (XCO) version 4.1.0 introduces an intelligent firmware upgrade auto-recovery mechanism designed to handle failed firmware downloads that occur after device activation and restart. This feature automatically recovers and completes failed firmware upgrades when network devices unexpectedly reboot during the update cycle.

During a firmware upgrade or restore operation, a device can reboot unexpectedly while maintenance mode is enabled. In earlier releases, the firmware download (FWDL) operation could remain in a failed state because the device temporarily lost connectivity.

With this enhancement, EFA automatically performs device recovery control (DRC) after the device reconnects and completes the remaining firmware workflow steps.

The firmware upgrade auto-recovery mechanism fulfills the following requirements:

- Automatic recovery from unexpected device reboots during firmware updates
- Reduced manual intervention requirements during firmware deployment
- Improved system resilience through intelligent state management
- Enhanced operational continuity during infrastructure maintenance
- Seamless recovery verification through system event logging
- Detects failed firmware download operations
- Verifies actual firmware version installed on the device
- Resumes the update process from the last successful step
- Completes remaining recovery steps automatically

Note

- Auto DRC starts only after a device update event is received.
- Device update events can occur through:
  - RASLog events
  - Manual device updates
  - Periodic device synchronization
- During recovery:
  - No additional firmware download events are published
  - The firmware download state tables are updated automatically
  - A firmware download completion event is published after successful recovery

Firmware Upgrade Workflow

The firmware upgrade workflow includes the following stages:

1. Download firmware
2. Enable maintenance mode (optional)
3. Copy running configuration to backup configuration
4. Activate firmware or restore image
5. Reload device
6. Poll for reload completion
7. Disable maintenance mode
8. Copy running configuration to backup configuration again
9. Commit firmware or complete restore

If the device reloads after the activation or restore stage, EFA resumes the workflow automatically from the `pollForReloadCompleted` stage.

Supported Scenarios

- Firmware Upgrade Failure After Activate or Restore Reload

  If the firmware upgrade or restore fails after the device reloads during activation or restore, EFA automatically resumes the workflow and updates the firmware status to Completed.

- Continue-on-Error Operations

  For firmware upgrades configured with `continue-on-error`, EFA performs recovery for all affected devices individually and completes the workflow.

Unsupported Scenarios

- Failure Before Activate Reload

  If the firmware workflow fails before the activate or restore reboot occurs, Auto DRC does not run. A new firmware download operation must be started manually.
- Fabric-Wide Stop-on-Error Operations

  If a fabric-wide firmware upgrade stops because of a `stop-on-error` condition, Auto DRC only resumes the failed device workflow. It does not continue the remaining fabric-wide upgrade operation automatically.

Firmware Upgrade Failure Scenarios

Learn about the firmware upgrade failure scenarios

Scenario Analysis

Scenario	Auto-Recovery	Expected Result	Notes
Pre-Activation Failure	No	Failed Status	Manual Restart Required
Fabric-Wide Stop-on-Error	Yes (Device Only)	Device Completed	Fabric Halted
Fabric-Wide Continue-on-Error	Yes (All Devices)	All Completed	Seamless Completion

Scenario 1: Early-Stage Failure (No Auto-Recovery)

When firmware download fails before device activation and reboot, the firmware has not yet been activated on the device. In this scenario:

- Auto-recovery does not apply
- Firmware download status remains in failed state
- Administrator must initiate a new firmware download from the beginning

Scenario 2: Fabric-Wide Upgrade with Device Failure

When a fabric-wide firmware upgrade encounters a device failure with stop-on-error policy:

- Failed device undergoes auto-recovery based on failure point
- Remaining devices in fabric-wide upgrade are NOT automatically resumed
- Administrator must explicitly resume fabric-wide upgrade after resolving failed device

Scenario 3: Fabric-Wide Upgrade with Continue-on-Error

When fabric-wide upgrade is configured with continue-on-error policy:

- Auto-recovery executes for all failed devices independently

- Remaining devices continue upgrade process without interruption
- Fabric-wide upgrade completes successfully across all devices

Firmware Upgrade Auto Recovery Flow

Automatic Detection and Initiation

Auto-recovery is triggered automatically through the device update process. When XCO detects a device status change (such as reconnection after a reboot), it performs intelligent checks:

- Identifies devices with firmware downloads in a failed state
- Verifies the actual firmware version currently running on the device
- Compares current version with target firmware version
- Automatically initiates recovery process if firmware matches expected version

Recovery Steps

Once auto-recovery is initiated, the system executes the following completion steps in sequence:

1. **Step 1: Maintenance Mode Disable**

   If the device was operating in maintenance mode during the update process, auto-recovery automatically disables maintenance mode to restore normal operational state. If maintenance mode was not enabled, this step is skipped.
2. **Step 2: Configuration Persistence**

   The system copies the active running configuration to persistent backup storage. This step ensures configuration consistency and can be executed multiple times without negative impact.
3. **Step 3: Firmware Commitment**

   If firmware commitment is configured, the system commits the activated firmware to permanent device storage, completing the update cycle. This step completes within seconds and is resilient to device operations.

System State Management

During auto-recovery, XCO maintains consistent system state through intelligent management:

- No "firmware download in progress" flag is set, preventing duplicate operations
- System state tables are updated with recovery progress
- Firmware download completion event is published upon success
- Operation history is maintained for audit and troubleshooting

Monitoring and Verification

- Monitor firmware upgrade status through XCO management interface
- Maintain records of firmware upgrade operations and outcomes
- Verify firmware version alignment across infrastructure

- Periodically review system event logs for auto-recovery execution
- Document any manual interventions required during upgrades
- Confirm the device reconnects successfully after reboot
- Verify maintenance mode is disabled automatically
- Confirm firmware download status changes to Completed
- Verify firmware commit operation completes successfully
- Ensure firmware completion events are generated correctly

Upgrade Strategy Selection

Choose appropriate error handling policies for fabric-wide upgrades:

- Stop-on-error: Preferred when maximum stability is required
- Continue-on-error: Suitable for stable environments with auto-recovery enabled
- Consider infrastructure size and criticality when selecting strategy

Device Health Management

Device Health Management (DHM) performs drift and reconciliation services, restoring fabric-related configurations. For RMA workflow, refer to the *Return Material Authorization* topic.

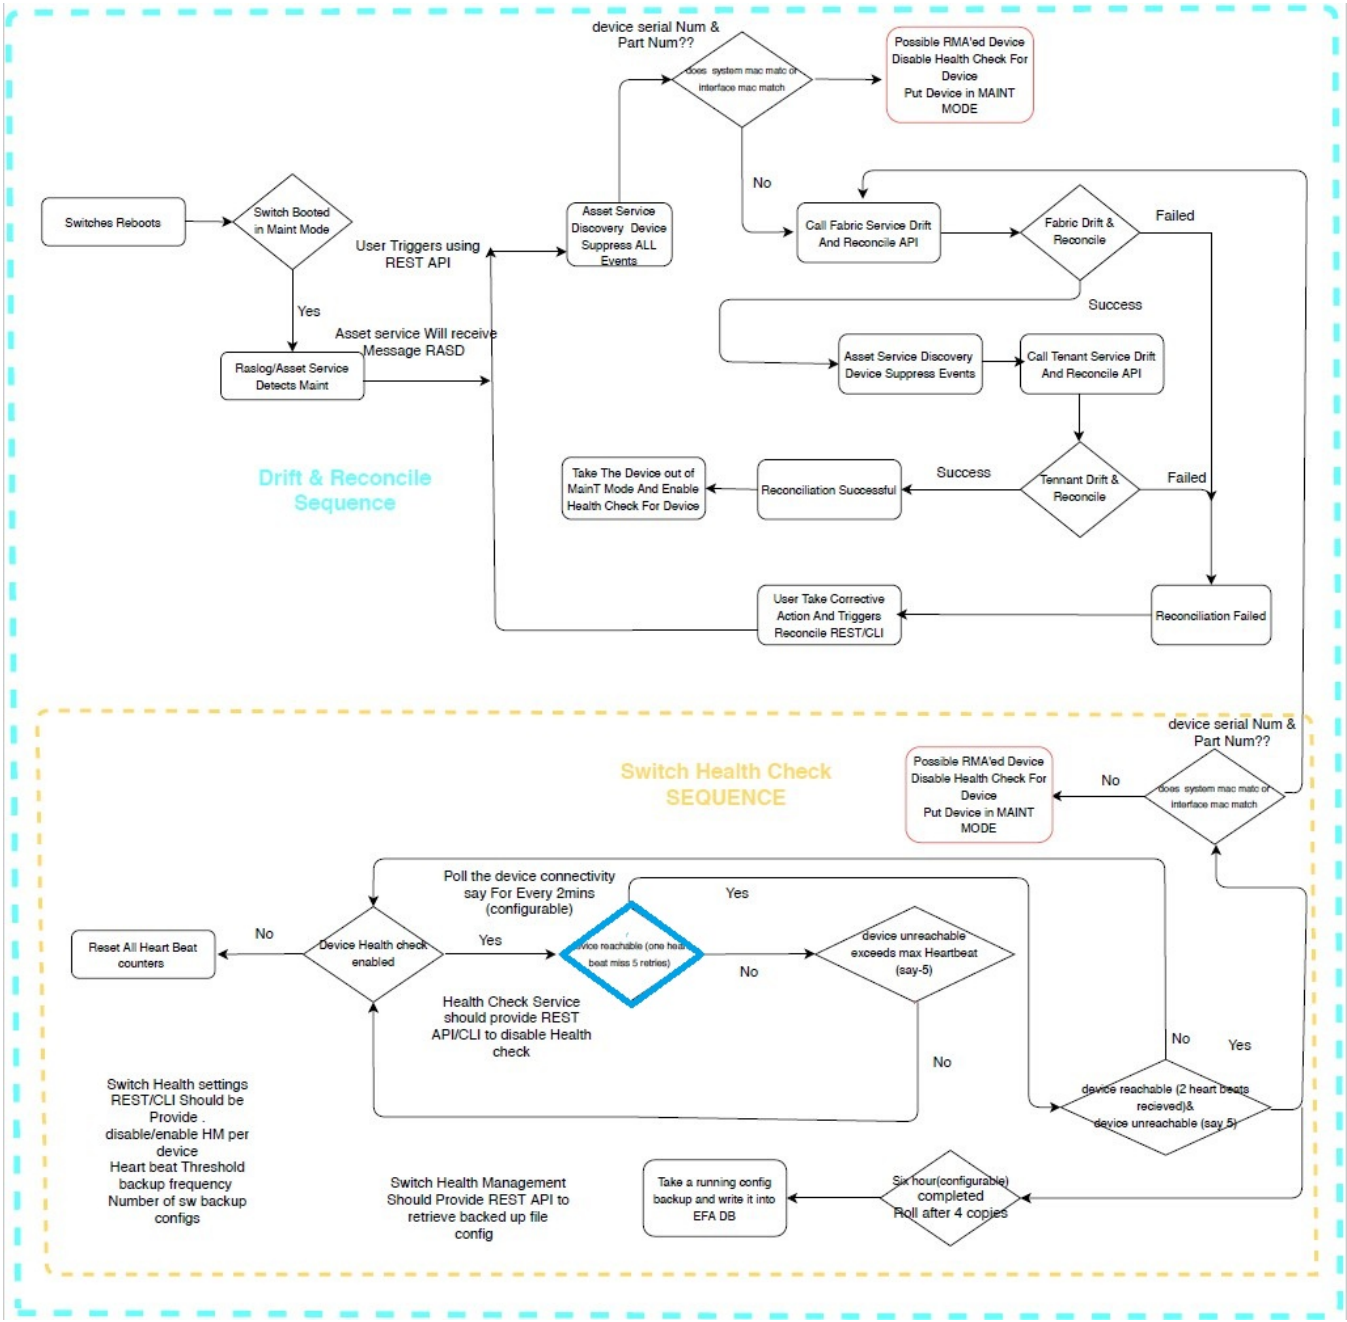

Figure 53: Device Health Management workflow

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.

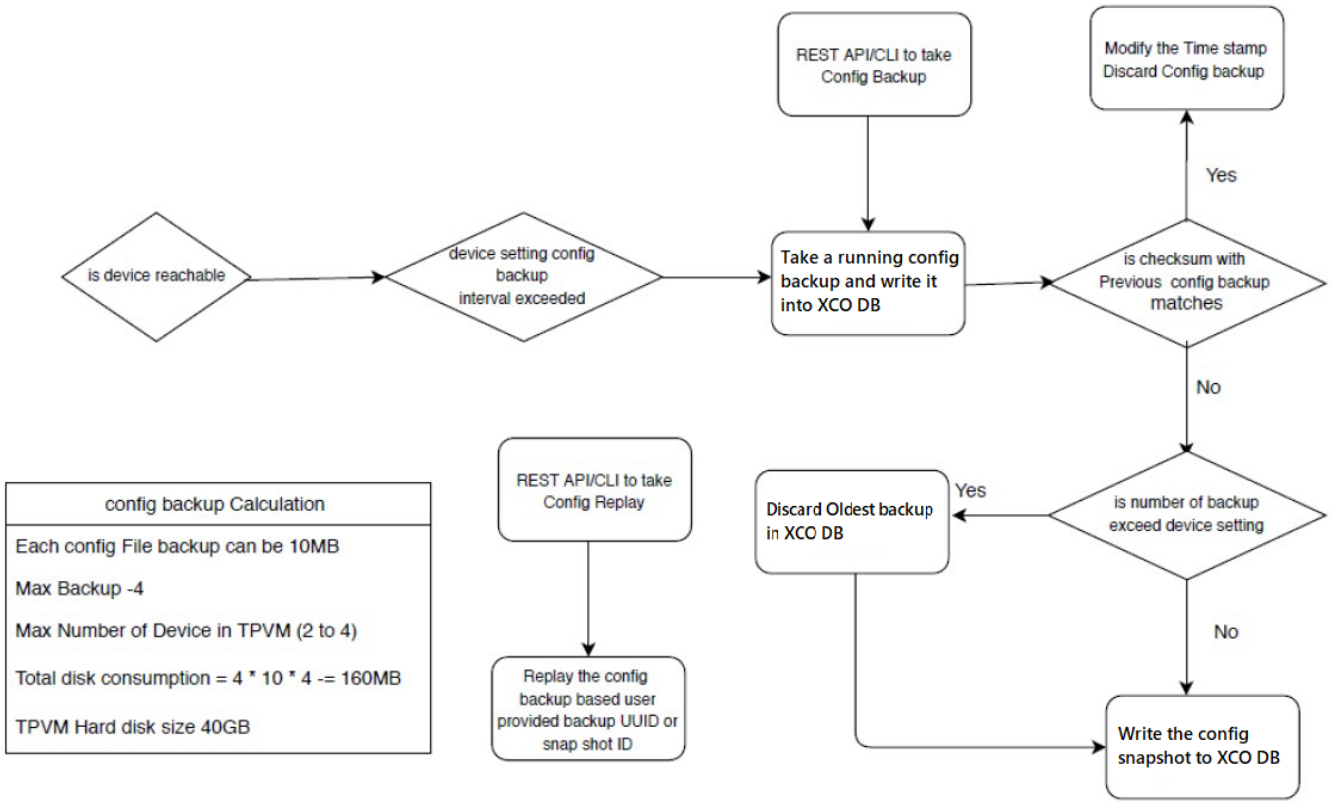

Figure 54: Workflow

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 on page 774.

   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 776. For Device Health Management workflow. refer to the *Device Health Management* topic.

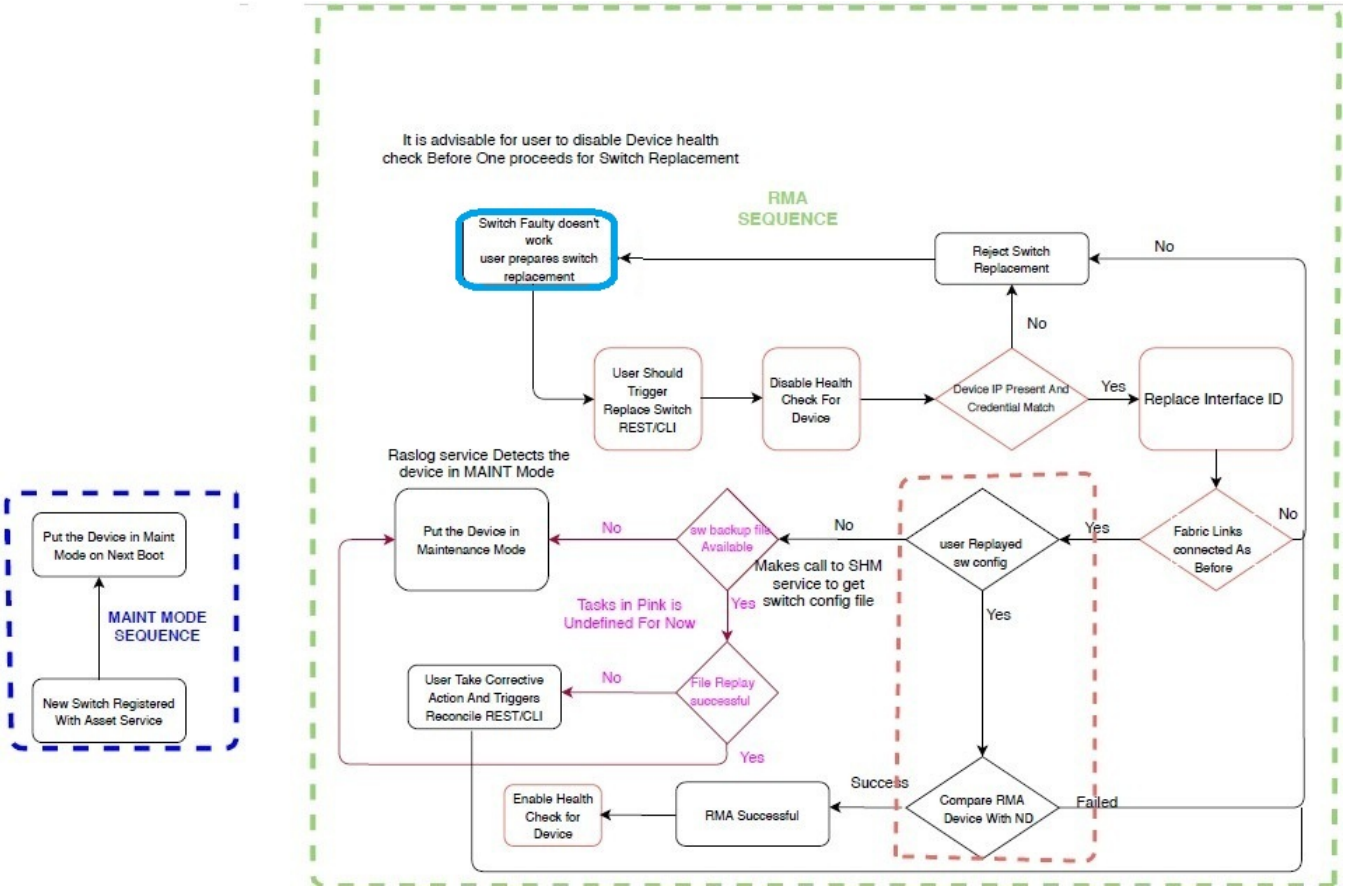

Figure 55: RMA workflow

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 125.

     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 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 775.

Procedure

1. Obtain the configuration backup of the old device.

   For more information, see Configure Backup and Replay on page 774.

   ```
   # 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 125.

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.

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.

Display Inventory Device Interface

You can display interfaces for each device.

About This Task

Follow this procedure to display and verify interfaces for each device.

You can verify that the interfaces are in the right state. The command helps you to see the current state and some details of the physical interfaces including breakout ports from XCO.

Procedure

Run the following command to display details of the physical interfaces:

```
efa inventory device interface list
```

This command displays the list of interfaces and details for the specified IP address, including the application state that indicates whether the device configuration is synchronized with XCO or has drifted (refreshed or deleted).

Example

The following example displays the physical interfaces of a device:

```
efa inventory device interface list --ip 10.20.63.140
+--------------+------+-----------+--------+-------+-------------+--------
+-------------------+------------+--------+-------------+
|   DeviceIP   | Name | Interface | Admin
| Oper  | Description | Line   |        MAC        | Switchport | IP     |  App State  |
|              |      | Type      | Status | Status|
| Speed  |                  | Mode       | Address|             |
+--------------+------+-----------+--------+-------+-------------+--------
+-------------------+------------+--------+-------------+
| 10.20.63.140 | 0/1  | ethernet  | up
| down  |              | unknown| 00:04:96:d6:fd:61 | unknown    |        | cfg-in-sync |
+              +------+-----------+--------+-------+-------------+--------
+-------------------+------------+--------+-------------+
|              | 0/2  | ethernet  | up
| down  |              | unknown| 00:04:96:d6:fd:62 | unknown    |        | cfg-in-sync |
+              +------+-----------+--------+-------+-------------+--------
+-------------------+------------+--------+-------------+
|              | 0/3  | ethernet  | up
| down  |              | unknown| 00:04:96:d6:fd:63 | unknown    |        | cfg-in-sync |
+              +------+-----------+--------+-------+-------------+--------
+-------------------+------------+--------+-------------+
|              | 0/4  | ethernet  | up
| down  |              | unknown| 00:04:96:d6:fd:64 | unknown    |        | cfg-in-sync |
+              +------+-----------+--------+-------+-------------+--------
+-------------------+------------+--------+-------------+
```

Display LLDP Inventory Device

You can display LLDP neighbors for each device.

About This Task

Follow this procedure to display LLDP neighbors for each device.

Procedure

To display details of LLDP neighbors, run the following command:

```
efa inventory device lldp list
```

The **`efa inventory device lldp list`** lists the LLDP (Link Layer Discovery Protocol) neighbors for a device.

Example

The following example displays LLDP neighbors of a device:

```
(efa:extreme)extreme@tpvm-111:~$ efa inventory device lldp  list  --ip 10.20.63.140
+--------------+--------------+----------------+-------+--------------+-----------------
+------------+----------------------+
|   DeviceIP   | Local        | Local Int MAC  | PO    | Remote       | Remote
Interface| Remote     | Remote Interface     |
|              | Interface    |                | Number| Interface    | MAC
| System Name| Description          |
+--------------+--------------+----------------+-------+--------------+-----------------
+------------+----------------------+
| 10.20.63.140 | Ethernet 0/27| 0004.96d6.fd7b | 64    | Ethernet 0/27| 0004.96d7.104d
| SW-141     | clusterPeerIntfMember|
+              +--------------+----------------+-------+--------------+-----------------
+------------+----------------------+
```

```
|              | Ethernet 0/28| 0004.96d6.fd7b | 64     | Ethernet 0/28| 0004.96d7.104d
| SW-141     | clusterPeerIntfMember|
+              +--------------+----------------+--------+--------------+-----------------
+------------+----------------------+
|              | Ethernet 0/29| 0004.96d6.fd7d |        | Ethernet 0/27| 0004.96d6.f611
| S110       | Eth 0/27             |
+              +--------------+----------------+--------+--------------+-----------------
+------------+----------------------+
|              | Ethernet 0/30| 0004.96d6.fd7e |        | Ethernet 0/27| 0004.96d6.fc57
| S39        | Eth 0/27             |
+              +--------------+----------------+--------+--------------+-----------------
+------------+----------------------+
```

Configure Physical Port Speed

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. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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 ---
```

Device Running Config Persist

Use the **`device running config-persist`** command to save the running-config to startup-config on the devices from XCO.

The following example shows the output of the running-config persist on a specific device:

```
(efa:root)root@ubuntu:~# efa inventory device running-config persist --ip 10.20.48.161
Persist Device(s) Running-Config[success]
+--------------+-----------+--------+---------+--------+
|  IP Address  | Host Name | Fabric | Status  | Reason |
+--------------+-----------+--------+---------+--------+
| 10.20.48.161 | leaf1     | fs     | Success |        |
+--------------+-----------+--------+---------+--------+
Persist Running-Config Details
--- Time Elapsed: 10.10378192s ---
```

The following example output shows that fabric option is used to perform running config persist on all fabric devices:

```
(efa:root)root@ubuntu:~# efa inventory device running-config persist --fabric fs
Persist Device(s) Running-Config[success]
+--------------+-----------+--------+---------+--------+
|  IP Address  | Host Name | Fabric | Status  | Reason |
+--------------+-----------+--------+---------+--------+
| 10.20.48.161 | leaf1     | fs     | Success |        |
+--------------+-----------+--------+---------+--------+
| 10.20.48.162 | leaf2     | fs     | Success |        |
+--------------+-----------+--------+---------+--------+
Persist Running-Config Details
--- Time Elapsed: 10.115473486s ---
```

Device Execute CLI

Device Execute CLI runs a specified command on one or more devices.

The following example shows the output of a device execute CLI which runs the **`show run int eth 0/5`** command:

```
(efa:root)root@ubuntu:~# efa inventory device execute-cli --command "show run int eth
0/5-6"  --ip 10.20.48.161

Execute CLI[success]
+--------------+-----------+--------+-------------------------+---------+--------
+--------------------------------+
|  IP Address  | Host Name | Fabric |        Command
| Status  | Reason |            Output             |
+--------------+-----------+--------+-------------------------+---------+--------
+--------------------------------+
| 10.20.48.161 | leaf1     | fs     | show run
int eth 0/5-6 | Success |        | leaf1# show run int eth 0/5-6 |
|              |           |        |                           |
|        | interface Ethernet 0/5        |
|              |           |        |                           |
|        |  no shutdown                  |
|              |           |        |                           |
|        | !                             |
```

```
|              |           |        |                          |
|        | interface Ethernet 0/6        |
|              |           |        |                          |
|        |  no shutdown                  |
|              |           |        |                          |
|        | !                             |
|              |           |        |                          |
|        |                               |
+--------------+-----------+--------+--------------------------+---------+--------
+--------------------------------+
Execute CLI Details
--- Time Elapsed: 6.27803233s ---
```

Using fabric option, you can use the execute CLI to run a specified command on all the devices of a fabric.

```
(efa:root)root@ubuntu:~# efa inventory device execute-cli --command "show run int eth
0/5"  --fabric fs
Execute CLI[success]
+--------------+-----------+--------+----------------------+---------+--------
+------------------------------+
|  IP Address  | Host Name | Fabric |       Command
| Status  | Reason |           Output            |
+--------------+-----------+--------+----------------------+---------+--------
+------------------------------+
| 10.20.48.161 | leaf1     | fs     | show run
int eth 0/5 | Success |        | leaf1# show run int eth 0/5 |
|              |           |        |                          |
|        | interface Ethernet 0/5      |
|              |           |        |                          |
|        |  no shutdown                 |
|              |           |        |                          |
|        | !                            |
|              |           |        |                          |
|        |                              |
+--------------+-----------+--------+----------------------+---------+--------
+------------------------------+
| 10.20.48.162 | leaf2     | fs     | show
run int eth 0/5 | Success |        | leaf2# show run int eth 0/5 |
|              |           |        |                          |
|        | interface Ethernet 0/5      |
|              |           |        |                          |
|        |  no shutdown                 |
|              |           |        |                          |
|        | !                            |
|              |           |        |                          |
|        |                              |
+--------------+-----------+--------+----------------------+---------+--------
+------------------------------+
Execute CLI Details
--- Time Elapsed: 5.132525111s ---
```

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. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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 ---
```

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 ---
```

3. To list the breakout interfaces from XCO, run the **`efa inventory device interface list-breakout`** command.

   The **`efa inventory device interface list-breakout`** lists breakout ports, including the application state that indicates if the configuration on the device is in sync or has drifted (refreshed or deleted) with respect to XCO.

   This example lists all breakout interfaces created on devices.

```
$ efa inventory device interface list-breakout --ip 10.20.246.18

+--------------+--------+-----------------+
|  IP Address  |  Name  |    AppState     |
+--------------+--------+-----------------+
| 10.20.246.18 | 0/52:1 | cfg-refreshed   |
+              +--------+-----------------+
|              | 0/52:2 | cfg-refreshed   |
+              +--------+-----------------+
|              | 0/52:3 | cfg-refreshed   |
+              +--------+-----------------+
|              | 0/52:4 | cfg-refreshed   |
+              +--------+-----------------+
|              | 0/53:1 | cfg-in-sync     |
+              +--------+-----------------+
|              | 0/53:2 | cfg-in-sync     |
+              +--------+-----------------+
|              | 0/53:3 | cfg-in-sync     |
+              +--------+-----------------+
```

```
|              | 0/53:4 | cfg-in-sync     |
+--------------+--------+-----------------+
```

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. For more configuration examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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 ---
```

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.

Interface Configuration Behavior Changes

Learn about end-user-visible behavior changes introduced by the interface configuration tracking update for XCO 4.1.0. These changes affect efa show-running-config, drift detection, and drift reconcile behavior.

This update changes how XCO tracks interface settings. XCO now distinguishes between values that a user explicitly configured through XCO and values that came from device defaults. This improves accuracy in running configuration output and drift handling.

What changed

The update introduces four main behavior changes:

- **Running configuration output is more accurate.** XCO keeps certain explicitly configured interface values in efa show-running-config even when they match default or automatic device behavior.
- **Drift detection is narrower.** XCO now detects drift only for fields that XCO explicitly configured.
- **Drift reconcile is safer.** XCO avoids restoring some interface settings unless XCO owns that configuration.
- **Some upgrade results are visible.** After upgrade, you may see additional commands in running configuration for settings that were already stored as user configuration.

Changes to Running Configuration Output

After this update, the **`efa show-running-config`** command can display interface commands that were previously omitted after a device refresh or update.

Scenario	Previous behavior	New behavior
Redundant management explicitly disabled	Only enabled state was shown	Disabled state is also shown as an explicit command
FEC set to automatic mode	The setting could disappear from running configuration	The setting remains visible if the user explicitly set it
Speed set to automatic mode	The setting could disappear from running configuration	The setting remains visible if the user explicitly set it
MTU set to the same value as the device global MTU	The setting could disappear from running configuration	The setting remains visible if the user explicitly set it

Examples

Redundant management disabled

If you explicitly disable redundant management through XCO , the command now remains visible in running configuration.

```
efa inventory device interface redundant-management --ip "10.64.188.117" --if-name 0/1
--enable false
```

FEC automatic mode

If you explicitly set FEC to automatic negotiation, XCO now tracks that choice as user configuration.

```
efa inventory device interface set-fec --ip "10.64.188.117" --if-type eth --if-name
0/1 --mode auto-negotiation
```

Speed automatic mode

If you explicitly set interface speed to automatic mode, XCO preserves that setting in running configuration after device updates.

MTU matches global MTU

If you explicitly configure an interface MTU and that value matches the device global MTU, XCO still shows the command in running configuration.

Note
If a value appears in efa show-running-config, it now more accurately represents an explicit user action via XCO, rather than merely a non-default device state.

Changes to Drift Detection

XCO now detects drift only for interface fields that XCO explicitly configured through EFA CLI or API.

Behavior area	Previous behavior	New behavior
Field ownership	Any non-empty or non-zero field could be treated as XCO-managed	Only explicitly configured XCO fields are treated as XCO-managed
Out-of-band device changes	Could trigger false drift	Do not trigger drift if XCO never configured that field
Drift scope	Broader and sometimes noisy	Narrower and more accurate

This means XCO no longer tries to protect or restore fields that were never configured through XCO.

Note

- If a network administrator changes a field directly on the device, and XCO never configured that field for the interface, XCO does not report drift for that change.
- When investigating drift, first confirm whether the field was originally configured through XCO. If XCO did not configure it, XCO does not own it for drift purposes.

Changes to Drift Reconcile Behavior

This update also improves drift reconcile behavior for redundant management.

Previously, XCO could attempt to restore redundant management when the stored value was enabled, even in cases where the port no longer supported that configuration. For example, this could happen after an out-of-band breakout change on a parent port.

Now, XCO restores redundant management only when both of these conditions are true:

- XCO explicitly configured the setting
- The stored value is enabled

This reduces unnecessary drift reconcile failures on unsupported ports.

Result: Drift reconcile is now safer because XCO restores only the interface settings that it explicitly owns.

FEC Automatic Mode is Now Fully Tracked

FEC automatic mode has a specific behavior improvement. Earlier, setting FEC to automatic mode could be lost from running configuration after later device updates. With this update, XCO treats FEC automatic mode as a complete tracked configuration.

XCO now does all of the following for FEC automatic mode:

- Stores the configuration
- Shows it in efa show-running-config

- Detects drift if the value changes out of band
- Restores it during drift reconcile

Upgrade impact

After upgrade to XCO 4.1.0, you may notice some changes in displayed running configuration.

- **Additional running-config entries may appear.** If existing interface settings were already stored as user-configured values, the migration can cause them to appear in efa show-running-config.
- **No expected drift surge.** The migration seeds the tracking data from existing configuration records, so false drift after upgrade is not expected.
- **Previously invisible explicit settings may now remain visible.** This is expected behavior.

Upgrade note After upgrade, seeing more interface commands in efa show-running-config does not necessarily indicate new configuration was pushed. It can reflect improved tracking of existing user-configured values.

Warning
After upgrade, seeing more interface commands in efa show-running-config does not necessarily indicate new configuration was pushed. It can reflect improved tracking of existing user-configured values.

What Users Should Expect

- Running configuration output may include more interface commands than before.
- Explicitly configured automatic values, such as FEC auto and speed auto, persist more consistently.
- Explicitly disabled redundant management is now visible.
- Drift alarms should be more accurate because XCO ignores fields it never configured.
- Drift reconcile operations should fail less often in interface ownership edge cases.

Quick reference

Area	User-visible outcome	Why it matters
`efa show-running-config`	Shows more explicit interface settings	Reflects user intent more accurately
Drift detection	Ignores unmanaged out-of-band changes	Reduces false drift
Drift reconcile	Restores only XCO-owned values in key cases	Reduces reconcile failures
FEC auto handling	Persists and is protected	Makes automatic mode a first-class configuration state

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, 4.1.0*.

Note

- When SLX is in maintenance mode, the SLX interface state changes do not sync with XCO. You must manually run the **`efa inventory device update --ip="SLX IP"`** command in XCO for XCO to reflect SLX's latest interface state.
- Starting from XCO 3.8.0, admin state settings for Ethernet interfaces will be exclusively handled by the inventory service, replacing tenant service.

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 ---
```

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 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
Error :  Please specify any one of Ip Address or Fabric Name.

update device setting

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 OS ONE device 30s-24h
Example. 30s or 99m or 1h20m or 20m, default 6m for Fabric device, 30s for OS ONE
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 8,16,32,64,128 and 0 to de-configure
    --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 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 Extreme 8720 and Extreme 8520 hardware
    --ip-option-disable string                       Enter Yes to disable processing
IP packets with IP options.  Enter No to enable processing IP packets with ip options
    --ip-option-disable-cpu string                   Enter Yes to disable processing
IP packets with IP options destined to the CPU. Enter No to enable processing IP
packets with IP options destined to the CPU
    --ipv6-option-disable string                     Enter Yes to disable processing
IPV6 packets with IP options.  Enter No to enable processing IPV6 packets with IP
options
  --- Time Elapsed: 54.634847ms ---
```

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     | 128    |
+--------------------------------+-------+
| Maximum Ipv6 Prefix Length 64  | Yes   |
+--------------------------------+-------+
| Urpf                           | Yes   |
+--------------------------------+-------+
--- Time Elapsed: 56.1149ms ---
```

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 Port Dampening on Interface

You can configure port dampening on SLX interface to minimize excessive interface flapping.

About This Task

Follow this procedure to configure port dampening on SLX interface.

Procedure

1. To configure port dampening on the SLX interface, run the **`set-link-error-disable`** command.

```
(efa:root)root@ubuntu:~# efa inventory device interface set-link-error-disable --ip
10.20.48.161  --if-name 0/20 --toggle-threshold 10 --sampling-time 20 --wait-time 10
+--------------+----+------+------------------+---------------+-----------+---------
+--------+
|   DeviceIP   | ID | Name | Toggle Threshold | Sampling Time | Wait
Time | Result  | Reason |
+--------------+----+------+------------------+---------------+-----------+---------
+--------+
| 10.20.48.161 | 11 | 0/20 | 10               | 20            | 10
| Success |        |
+--------------+----+------+------------------+---------------+-----------+---------
+--------+
Interface Details
--- Time Elapsed: 5.982855992s ---
```

2. Complete the following configuration on SLX devices:

```
leaf1# sh run int eth 0/20
interface Ethernet 0/20
 link-error-disable 10 20 10
 no shutdown
!
```

Configure Description on Device Interface

You can configure the description on the device interface. You can also create custom descriptions for Ethernet ports and link aggregation groups on SLX devices.

About This Task

Note
Starting from XCO 3.8.0, the behavior of Fabric and Tenant services regarding interface descriptions will change:

- Fabric Service
  1. After configuring fabric, the descriptions of MCT interfaces will no longer be changed to "clusterPeerIntfMember", and fabric port descriptions will no longer be changed to "Link to <ip> <role>"
  2. After deleting fabric, existing interface descriptions will remain unchanged, if it is already set.
- Tenant Service
  1. After creating a tenant PO, interface descriptions will no longer be changed to "Port-channel <name> Member interface".
  2. After deleting a Tenant PO, existing interface descriptions will remain unchanged, it is already set
- Starting from XCO 3.8.0, description management for Ethernet interfaces will be exclusively handled by the inventory service, replacing fabric and tenant services.

- Drift and reconcile (DRC): The description settings configured using the **`set-description`** command will be managed by the Inventory Service. Only the Inventory Service will perform drift and reconciliation for these settings. As a result, Fabric and Tenant services will not display any drift status for descriptions, since they are not responsible for setting them.
- XCO Upgrade: Existing description configurations on switch interfaces will remain unchanged during the upgrade process. However, the Fabric and Tenant services will not detect any drift for the description attribute after the upgrade.

Follow this procedure to configure (set and unset) the description on the device interface.

Note
For more information on syntax and command examples, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

1. Set the Description.
   a. Using CLI: At the interface level, run the **`efa inventory device interface set-description`** command to set the description.
   The following example sets the description on the device interface:

```
$ efa inventory device interface set-description -ip 10.64.164.47 -if-name 0/14 --
description "interface eth 0/14"

+--------------+----+------+--------------+--------------------+---------+-------+
|   DeviceIP   | ID | Name | Interface Type| Description        | Result  | Reason|
+--------------+----+------+--------------+--------------------+-- ------+-------+
```

```
| 10.64.164.47 | 95 | 0/14 | ethernet      | interface eth 0/14 | Success |       |
+--------------+----+------+---------------+--------------------+---------+-------+
```

b. Using REST API: Use the following REST API to set the description.

```
PUT  " http://goinventory-service:8082/v1/inventory/interfaces/description"

 curl --request PUT "http://goinventory-service:8082/v1/inventory/
interfaces/description" -d '{"ip_address":["10.64.196.52"], "intfNames":
["0/32"],"intfDescription":"interface ethernet"}'

[{"ip_address":"10.64.196.52","interfaces":
[{"id":63,"intf_type":"ethernet","name":"0/32","description":"interface
ethernet","status":{"result":"Success"}}]}]
```

2. Unset the description.

a. Using CLI: At the interface level, run the **`efa inventory device interface unset-description`** command to set the description.

```
$ efa inventory device interface unset-description --ip 10.64.212.11 --if-name
"0/37"
+--------------+------+------+----------------+---------+--------+
|   DeviceIP   |  ID  | Name | Interface Type | Result  | Reason |
+--------------+------+------+----------------+---------+--------+
| 10.64.212.11 | 2265 | 0/37 | ethernet       | Success |        |
+--------------+------+------+----------------+---------+--------+
Interface Description Details
--- Time Elapsed: 5.383109148s ---
```

b. Using REST API: Use the following REST API to unset the description:

```
DELETE "http://goinventory-service:8082/v1/inventory/interfaces/description"

curl --request DELETE "http://goinventory-service:8082/v1/inventory/interfaces/
description" -d '{"ip_address":["10.64.196.52"], "intfNames":["0/32"]}'

[{"ip_address":"10.64.196.52","interfaces":
[{"id":63,"intf_type":"ethernet","name":"0/32","status":{"result":"Success"}}]}]
```

3. At the interface level, run the **`efa inventory device interface list`** command to fetch the device inventory list.

```
$ efa inventory device interface list --ip 10.64.164.47
+--------------+------+---------------+--------------+-------------+-----------------------
+------------+------------------+----------------+------------+-------------+
|   DeviceIP   | Name | Interface Type | Admin Status | Oper Status |
Description      | Line Speed |        MAC        | Switchport Mode | IP Address |  App State  |
+--------------+------+---------------+--------------+-------------+-----------------------
+------------+------------------+----------------+------------+-------------+
| 10.64.164.47 | 0/1  | ethernet       | up           | down        |
| unknown    | 00:04:96:b8:ce:0e | unknown         |            | cfg-in-sync |
+              +------+----------------+--------------+-------------+-----------------------
+------------+------------------+----------------+------------+-------------+
|              | 0/2  | ethernet       | up           | down        |
| unknown    | 00:04:96:b8:ce:0f | unknown         |            | cfg-in-sync |
+              +------+----------------+--------------+-------------+-----------------------
+------------+------------------+----------------+------------+-------------+
|              | 0/3  | ethernet       | up           | down        |
| unknown    | 00:04:96:b8:ce:10 | unknown         |            | cfg-in-sync |
+              +------+----------------+--------------+-------------+-----------------------
+------------+------------------+----------------+------------+-------------+
|              | 0/4  | ethernet       | up           | down        |
| unknown    | 00:04:96:b8:ce:11 | unknown         |            | cfg-in-sync |
+              +------+----------------+--------------+-------------+-----------------------
+------------+------------------+----------------+------------+-------------+
|              | 0/5  | ethernet       | up           | up          |
```

```
| 100Gbps    | 00:04:96:b8:ce:12 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/6  | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:13 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/7  | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:14 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/8  | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:15 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/9  | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:16 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/10 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:17 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/11 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:18 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/12 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:19 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/13 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:1a | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/14 | ethernet       | up          | down       | interface
eth 0/14    | unknown    | 00:04:96:b8:ce:1b | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/15 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:1c | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/16 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:1d | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/17 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:1e | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/18 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:1f | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/19 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:20 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/20 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:21 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/21 | ethernet       | up          | down       |
```

```
| unknown    | 00:04:96:b8:ce:22 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/22 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:23 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/23 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:24 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/24 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:25 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/25 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:26 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/26 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:27 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/27 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:28 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/28 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:29 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/29 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:2a | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/30 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:2b | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/31 | ethernet       | up          | up         | clusterPeerIntfMember
| 100Gbps    | 00:04:96:b8:ce:2c | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/32 | ethernet       | up          | up         | clusterPeerIntfMember
| 100Gbps    | 00:04:96:b8:ce:2d | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/33 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:2e | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/34 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:2f | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/35 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:30 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/36 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:31 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------+-----------------------
+------------+------------------+----------------+-------------+------------+-------------+
|              | 0/37 | ethernet       | up          | down       |
```

```
| unknown    | 00:04:96:b8:ce:32 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+
|              | 0/38 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:33 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+
|              | 0/39 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:34 | unknown         |            | cfg-in-sync |
+              +------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+
|              | 0/40 | ethernet       | up          | down       |
| unknown    | 00:04:96:b8:ce:35 | unknown         |            | cfg-in-sync |
+--------------+------+----------------+-------------+------------+-------------------------
+------------+------------------+----------------+-------------+------------+
Interface Details
```

4. Complete the following configuration on SLX device:

```
# show running-config interface Ethernet 0/14
   interface Ethernet 0/14
   description interface eth 0/14
   no shutdown
!
```

Configure RME on SLX 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 SLX 9740, Extreme 8720, and Extreme 8520 platforms. The configuration you set is stored in the XCO database. DRC is supported.

XCO automatically sets the PPS (packets per second) value after RME is enabled. For Extreme 8720, and Extreme 8520 devices, PPS is set to 8000.

On SLX 9740, the BPS (bits per second) is set to 20000 Kbps after RME is enabled.

Note

For information about the commands on RME, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*

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
!
```

Interface FEC

You can configure Forward Error Correction (FEC) on the SLX interface.

About This Task

Follow this procedure to configure Forward Error Correction (FEC) on the SLX interface.

The configuration set by you will continue to exist in the XCO database. DRC is supported. The default value of FEC configured by SLX is auto (auto-negotiation).

Procedure

1. Run the **`efa inventory device interface set-fec`** command to configure Forward Error Correction (FEC) on the SLX interface.

```
(efa:root)root@ubuntu:~# efa inventory device interface set-fec --ip 10.20.48.161 --if-
type eth --if-name 0/20 --mode rs-fec
+--------------+----+------+----------------+----------+---------+--------+
|   DeviceIP   | ID | Name | Interface Type | FEC Mode | Result  | Reason |
+--------------+----+------+----------------+----------+---------+--------+
| 10.20.48.161 | 9  | 0/20 | ethernet       | rs-fec   | Success |        |
+--------------+----+------+----------------+----------+---------+--------+
Interface Details
--- Time Elapsed: 4.934749593s ---
```

2. Complete the following configuration on SLX devices:

```
leaf1# sh run int eth 0/20
 interface Ethernet 0/20
 fec mode RS-FEC
 no shutdown !
```

3. Run the **`efa inventory device interface unset-fec`** command to unset FEC on the interface:

```
(efa:root)root@ubuntu:~# efa inventory device interface unset-fec --ip 10.20.48.161 --
if-type eth --if-name 0/20
+--------------+----+------+----------------+---------+--------+
| DeviceIP     | ID | Name | Interface Type | Result  | Reason |
+--------------+----+------+----------------+---------+--------+
| 10.20.48.161 | 0  | 0/20 | ethernet       | Success |        |
+--------------+----+------+----------------+---------+--------+
Interface Details
--- Time Elapsed: 4.906542529s ---
```

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.

XCO Native Support for SLX Threshold Monitor Settings

Set Threshold Monitor Options

You can configure the threshold monitor options.

About This Task
Follow this procedure to configure threshold monitor options.

Table 21: Drift Reconcile & Idempotency Support Table

Identify Drift	Reconcile configuration	Idempotency
Yes	Yes	Yes

For information about commands and supported parameters to configure threshold monitor options, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

Run the following command to set the threshold monitor options:

```
efa inventory device threshold-monitor set
```

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Example

- The following example sets a CPU threshold-monitor to trigger an SNMP event when the CPU threshold of 80% is exceeded. Ensure that the threshold must be exceeded three times (retry) with a polling interval of 60 seconds for the event to be triggered. This means that it would take three 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 --retry 3
  ```

- The following example sets a mac-table threshold-monitor to trigger an SNMP event when either the high-limit threshold of 80% or the low-limit threshold of 50% is exceeded. Events can be generated for the threshold exceeding the 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
  ```

Remove Settings for Threshold Monitor Options

You can reset all of the threshold monitor options back to their default values by removing the existing settings.

Procedure

Run the following command to remove existing settings for the threshold monitor options:

```
efa inventory device threshold-monitor unset
```

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Example

- The following example removes all the current inventory threshold-monitor settings for the specified device:
  ```
  efa inventory device threshold-monitor unset --ip 10.10.10.75-76
  ```
- The following example removes 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 view threshold monitor settings.

About This Task

Follow this procedure to show threshold monitor settings.

Procedure

Run the following command to display the threshold monitor settings:

```
efa inventory device threshold-monitor list
```

For more information on syntax and command examples, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Example

The following example shows the current inventory threshold-monitor settings for the specified device:

```
efa inventory device threshold-monitor list --ip 10.10.10.75
+
Basic Monitoring Type Configuration
+--------------+--------------+---------+------------+-----------+-------------+
|  IP Address  |     Type     | Actions | High Limit | Low Limit |  AppState   |
+--------------+--------------+---------+------------+-----------+-------------+
| 10.10.10.75 | mac-table    | raslog  | 100        | 70        | cfg-in-sync |
+              +--------------+---------+------------+-----------+-------------+
|              | route        | raslog  | 90         | 70        | cfg-in-sync |
+              +--------------+---------+------------+-----------+-------------+
|              | host         | raslog  | 100        | 70        | cfg-in-sync |
+              +--------------+---------+------------+-----------+-------------+
|              | lif          | raslog  | 100        | 70        | cfg-in-sync |
+              +--------------+---------+------------+-----------+-------------+
|              | nexthop      | raslog  | 100        | 70        | cfg-in-sync |
+              +--------------+---------+------------+-----------+-------------+
```

```
|              | ecmp         | raslog  | 100        | 70         | cfg-in-sync |
+              +--------------+---------+------------+------------+-------------+
|              | bfd-session  | raslog  | 100        | 70         | cfg-in-sync |
+              +--------------+---------+------------+------------+-------------+
|              | vxlan-tunnel | raslog  | 100        | 70         | cfg-in-sync |
+--------------+--------------+---------+------------+------------+-------------+

CPU/Memory Type Configuration
+--------------+--------+---------+------------+-----------+----------+-------
+-------------+
|  IP Address  |  Type  | Actions | High Limit
| Low Limit | Interval | Retry |  AppState   |
+--------------+--------+---------+------------+-----------+----------+-------
+-------------+
| 10.10.10.75 | cpu    | raslog  | 100
| 10        | 3500     | 10    | cfg-in-sync |
+              +--------+---------+------------+-----------+----------+-------
+-------------+
|              | memory | raslog  | 100        |
10        | 3600     | 100   | cfg-in-sync |
+--------------+--------+---------+------------+-----------+----------+-------
+-------------+
Threshold Monitor details
--- Time Elapsed: 81.027865ms ---
```

Additional Threshold Monitor Types

SLX has a feature that sets thresholds, which triggers action events when the limits are reached. This feature was initially designed for SLX-OS 20.4.2 and above, and was included in the release EFA 3.1.0. The original threshold types were CPU, memory, bfd-session, lif, mac-table, and vxlan-tunnel.

SLXOS 20.5.3/XCO 3.5.0 and later includes additional threshold monitor entities or types such as route, host, nexthop, and ECMP. Moreover, there is a new hardware-resources type that holds common or shared configuration parameters for threshold types, excluding the CPU and memory types.

Re-worked Threshold Monitor Parameters

The threshold monitor parameters in EFA 3.1.0 were designed to be consistent with SLX-OS 20.4.2. In EFA 3.1.0, the CLI included parameters count and interval for each threshold type. However, this has changed for SLXOS 20.5.3/XCO 3.5.0 and later, where they are now common or shared values for all the threshold monitor types, except for CPU and memory, which can be found under the option hardware-resources.

Under hardware-resources, the count default value has remained at 4 but the default interval value has changed from 120 to 30. The high and low limit parameters remain under the individual types, with their respective defaults remaining unchanged.

The retry parameter remains the same and is still only applicable to the types CPU and memory. The interval parameter under memory and CPU is still separate from the hardware-resources value and remains the same.

CLI Migration

In XCO 3.5.0 and later, the threshold monitoring parameters count and interval for bfd-session, lif, mac-table, and vxlan-tunnel have a common or shared value.

This means that the individual count and interval values for these types will no longer be included. As a result, these settings will not be migrated from previous XCO releases. The common or shared values will be set to their default values.

Disable IP Option

You can disable IP options. The **ip option disable** command blocks packets that have IP options. The **ip option disable-cpu** command configures dropping of packets with the destination as the device's CPU (Control Plane Processing Unit).

About This Task

Follow this procedure to disable or enable processing of IP (IPv4 and IPv6) packets with IP options.

You can configure IPv4 and IPv6 options on a device.

- SLX 9540 and SLX 9640 do not support **ip option disable-cpu**.
- SLX 20.3.4 supports **ip option disable** and **ipv6 option disable**.
- SLX 20.4.2 supports **ip option disable-cpu**.

Note
For information about commands and supported parameters to enable or disable IPv4 and IPv6 options, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

Run the following command to update IPv4 and IPv6 options:

```
efa inventory device setting update --ip <device ips> --ip-option-disable { Yes | No }
--ip-option-disable-cpu  { Yes | No } --ipv6-option-disable { Yes | No }
```

Example

- The following example disables IP option processing and IPv6 option processing while allowing IP option processing destined to the CPU.

```
$ efa inventory device setting update --ip 10.10.10.1 --ip-option-disable yes --ip-
option-disable-cpu no --ipv6-option-disable yes
+-------------+----------------------+---------+-------+-------+
| IP ADDRESS  |         NAME         | STATUS  | VALUE | ERROR |
+-------------+----------------------+---------+-------+-------+
| 10.24.2.120 | ip-option-disable    | Success | Yes   |       |
+             +----------------------+---------+-------+-------+
|             | ip-option-disable-cpu | Success | No    |       |
+             +----------------------+---------+-------+-------+
|             | ipv6-option-disable  | Success | Yes   |       |
+-------------+----------------------+---------+-------+-------+
```

- The following example shows the current inventory device settings for a specified device when IP option disable is set to “Yes” and IPv6 option disable is set to “Yes”.

```
$ efa inventory device setting show --ip 10.10.10.1
+--------------------------------+-------+-------------+
|              NAME              | VALUE |  APP STATE  |
+--------------------------------+-------+-------------+
| Maintenance Mode Enable On     | No    |             |
| Reboot                         |       |             |
+--------------------------------+-------+-------------+
| Maintenance Mode Enable        | No    |             |
+--------------------------------+-------+-------------+
| Maintenance Convergence Time   | 100   | cfg-in-sync |
+--------------------------------+-------+-------------+
| 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    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Static Prefix Independent      | No    | cfg-in-sync |
| Convergence                    |       |             |
+--------------------------------+-------+-------------+
| Maximum Load Sharing Paths     | 128    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Maximum Ipv6 Prefix Length 64  | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
| Urpf                           | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
| IP Option Disable              | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
| IP Option Disable CPU          | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
| IPV6 Option Disable            | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
```

The following is an associated SLX configuration:

```
NHLeaf1# show running-config ip option
ip option disable
NHLeaf1# show running-config ipv6 option
ipv6 option disable
```

- The following example shows the current inventory device settings for a specified device when IP option disable-cpu is set to “Yes”.

```
$ efa inventory device setting show --ip 10.10.10.1
+--------------------------------+-------+-------------+
|              NAME              | VALUE |  APP STATE  |
+--------------------------------+-------+-------------+
| Maintenance Mode Enable On     | No    |             |
| Reboot                         |       |             |
+--------------------------------+-------+-------------+
| Maintenance Mode Enable        | No    |             |
```

```
+--------------------------------+-------+-------------+
| Maintenance Convergence Time   | 100   | cfg-in-sync |
+--------------------------------+-------+-------------+
| 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    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Static Prefix Independent      | No    | cfg-in-sync |
| Convergence                    |       |             |
+--------------------------------+-------+-------------+
| Maximum Load Sharing Paths     | 128    | cfg-in-sync |
+--------------------------------+-------+-------------+
| Maximum Ipv6 Prefix Length 64  | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
| Urpf                           | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
| IP Option Disable              | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
| IP Option Disable CPU          | Yes   | cfg-in-sync |
+--------------------------------+-------+-------------+
| IPV6 Option Disable            | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
```

The following is an associated SLX configuration:

```
NHLeaf1# show running-config ip option
ip option disable-cpu
NHLeaf1# show running-config ipv6 option
```

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.
- The maximum supported password length is 40 characters.

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, 4.1.0*.

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, 4.1.0*.

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, 4.1.0*.

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.

Passwordless SSH or SCP Support for Secure and Efficient Backup and Supportsave Transfers

You can enhance credential management security with SSH certificate-based authentication. It's a stronger alternative to traditional SSH key-based or password-based methods, protecting against unauthorized access and vulnerabilities. This approach streamlines remote server communications, perfect for critical tasks like backup storage and supportsave retrieval. It's scalable, manageable, and robust for safeguarding sensitive data.

For passwordless authentication, SSH keygen is often preferred for its direct integration with the SSH protocol, ease of use, and simplicity.

Secure Copy Protocol (SCP) credential management poses security concerns. Migrating to SSH certificate-based authentication enhances security, scalability, and manageability. SSH certificate-based authentication provides the followng benefits:

1. Improved protection against unauthorized access
2. Reduced vulnerabilities associated with traditional SSH key-based or password-based authentication
3. Enhanced security for critical tasks like backup storage and support save retrieval

The private key file must be present on all XCO nodes within the `efadata` directory. Specifically, for server deployments, the file must be located in the `/opt/efadata/certs` directory, while for TPVM setups, it must reside in the `/apps/efadata/certs` directory.

Password Input	Keys Input	XCO Supportsave and Backup	SLX Supportsave
Yes	Yes	Yes	Yes
No	Yes	Yes	No
Yes	No	Yes	Yes
No	No	No	No

OpenSSH Certificate-based Authentication

OpenSSH certificate-based authentication enhances security and scalability by leveraging digital certificates signed by a trusted Certificate Authority (CA).

Explore this topic to understand the essential components of OpenSSH Certificate-based Authentication and their functions.

Certificate Authority (CA)

A Certificate Authority consists of a key pair (public and private) used to sign other keys.

- The **CA's private key** signs certificates.
- The **CA's public key** is used to verify these certificates.
- These verify servers to clients.
- A host certificate is a server's public key signed by the CA, ensuring clients connect to the correct server.

Organizations like **XCO** can act as the CA, or users can opt for another OpenSSH host to serve as the CA.

Host Certificates

- These verify servers to clients.
- A host certificate is a server's public key signed by the CA, ensuring clients connect to the correct server.

User Certificates

- These verify users to servers.

- A user certificate is a user's public key signed by the CA, allowing users to log in without distributing public keys to every server.

Note
User certificates are often referred to as client certificates. In this context, XCO functions as the OpenSSH client.

Configure OpenSSH Certificate-based Authentication

You can configure OpenSSH certificate-based authentication.

About This Task
Follow this procedure to configure OpenSSH certificate-based authentication.

Procedure

1. Create a CA.
   a. Generate a key pair for the CA using `ssh-keygen`.
   b. Keep the private key secure and distribute the public key to servers and clients.
2. Sign the keys.
   a. Generate host and user keys (by the client), then sign them with the CA.
   b. This process attaches the CA's signature to the keys, creating certificates.
3. Configure OpenSSH servers.
   a. Add the CA's public key to the `TrustedUserCAKeys` directive in the server's `sshd_config` file to trust the CA.
   b. Host certificates are also added to the server's SSH configuration.
4. Configure OpenSSH clients.
   a. Clients, such as XCO, add the CA's public key to the `known_hosts` file to trust the CA.

      XCO acts as an OpenSSH client to facilitate file uploads to the remote server.
   b. Specify user certificates in the SSH client configuration for secure access to remote servers.

Configure Remote Settings

You can configure remote server settings.

About This Task

Follow this procedure to configure remote server settings using CLI and API.

Note

- A password can be used along with keys to collect the device supportsave. The XCO supportsave process first attempts using the passwordless feature and falls back to using the password if the initial attempt fails.
- To avoid using the current remote settings, reset them before applying updates.
- Ensure that the certificate paths specified for the CLI are accessible on the active node in a multi-node environment.
- Verify that the certificates' validity and expiration are consistent with the signed timezone.
- For more information on syntax and command examples, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

1. To configure the remote server settings using CLI, run the following command:
   a. With Password
   ```
   efa system settings update --remote-server-ip 10.32.85.5 --remote-server-username
   root --remote-server-directory /home/user/Downloads --remote-password password
   ```
   b. Without Password (SSH Certificate Based Passwordless)
   ```
   efa system settings update -remote-server-password password --remote-server-
   ip 10.32.85.5 --remote-server-username root --remote-server-directory /home/user/
   Downloads --certificate-key /home/user/id_rsa-cert.pub --private-key id_rsa --
   passphrase pkeyabc  --ca-publickey /home/user/ca.pub
   ```
2. To configure the remote server settings using REST API, run the following command:

```
PrivateKey: Filename should be passed as the value

Example: id_rsa

Passphrase: Passphrase should be passed as the string value

Example: pwdabc

CertificateKey: Content of the certificate should be passed as an input

Example: "ssh-rsa-cert-v01@openssh.com AAAAHHNzaC1yc2Et....................."

CAPublicKey: content of the ca public key should be passed as an input

Example: "ssh-rsa FAb3BlbnNzaC5jb20AAAAg........."


curl --location --request POST 'http://gosystem-service:80/v1/system/settings' \

--header 'Content-Type: application/json' \

--data-raw '{
```

```
"keyval": [
  {
    "value": "5",
    "key": "MaxBackupFiles"
  },
  {
    "value": "0 0 * *",
    "key": "BackupSchedule"
  },
  {
    "value": "5",
    "key": "MaxSsFiles"
  },
  {
    "value": "10.10.10.10 / 2000::1",
    "key": "RemoteServerIP"
  },
  {
    "value": "scp / ftp",
    "key": "RemoteTransferProtocol"
  },
  {
    "value": "username",
    "key": "RemoteServerUsername"
  },
  {
    "value": "password",
    "key": "RemoteServerPassword"
  },
  {
    "value": "/root/test",
    "key": "RemoteServerDirectory"
```

```
    },
    {
      "value": "Enabled",
      "key": "PeriodicDeviceConfigBackup"
    },
{
      "value": "id_rsa",
      "key": "PrivateKey"
    },
{
"value": pwdabc
"key": Passphrase
},
{
      "value": "ssh-rsa-cert-v01@openssh.com
AAAAHHNzaC1yc2EtY2VydC12MDFAb3BlbnNzaC5jb20AAAAg",
      "key": "CertificateKey"
    },
{
      "value": "ssh-rsa AAAAHHNzaC1yc2EtY2VydC12MDFAb3BlbnNzaC5jb20AAAAg",
      "key": "CAPublicKey"
    },

  ]
}
'
```

Configure Passwordless Authentication via SSH Certificate-Based Authentication

You can configure passwordless authentication using OpenSSH certificates as an authentication option.

About This Task

Follow this procedure to configure passwordless authentication via SSH certificate-based authentication.

Procedure

1. Complete the client (XCO) configuration.

   Use the private key and signed certificate pub key in the XCO CLI or REST input.

   a. Run the following command to generate the RSA key pair (public or private key):
   ```
   ssh-keygen -t rsa
   ```
   The **id_rsa** and **id_rsa.pub** files are generated.

   b. Copy the **ca.pub** key to the known hosts with **@cert-authority** as prefix. The **ca.pub** file is generated in Step 2 as the public key of the Certificate Authority.

   c. Get the generated **id_rsa-cert.pub** from the CA and use it for the SSH connection.

   Note

   - The client (user or host) seeking SSH server authentication presents a Certificate Authority (CA)-signed certificate instead of a raw public key.
   - The client holds a public or private key pair, with the public key embedded in a CA-signed certificate.
   - The client safeguards the private key and uses it during the SSH handshake to verify public key ownership.

2. Complete the certificate authority (CA) server configuration.

   Generate a CA certificate and use it for signing the client public key.

   a. Run the following command to generate the CA certificate:
   ```
   ssh-keygen -t rsa -f ca
   ```
   b. The CA signs the **id_rsa.pub** file. Copy the generated id_rsa-cert.pub file to the client **ssh-keygen -s ca -I root -n root -V +52w id_rsa.pub**. The **id_rsa.pub** key was generated by client (XCO)

   c. The CA signs the host public key of the remote sever and copy the generated **cert.pub** file to the remote server **ssh-keygen -s ca -I root -n root -V +52w -h sss_host_rsa_key.pub**.

   Note

   - The CA is a trusted entity signing client and server certificates.
   - The CA server holds a private key for signing certificates, asserting trust in the public key for specified users or hosts.
   - Both SSH clients and servers trust the CA for identity verification.

3. Complete the remote server configuration.

   Copy the CA public key to the remote server and configure the path in SSH.

   a. Copy **ca.pub** to the remote server and configure the path in the **sshd_config**.

   b. Get the signed host certificate and configure the path in the **sshd_config**.

   For example, **TrustedUserCAKeys /etc/ssh/ca.pub** and **HostCertificate ssh_host_rsa_key-cert.pub**.

c. Restart the SSH service after the config changes.

Note
The SSH server validates client certificates using the pre-configured CA public key. Certificate validation ensures the certificate:

- Was signed by the trusted CA.
- Is within the allowed time frame.
- Hasn't been revoked.
- Meets CA-set conditions.

Example
The following example output shows transfer of supportsave and backup to the remote server.

- **Supportsave**

  To capture a supportsave of the XCO and copy it to a remote server (after configuring remote server settings):

```
(efa:user)root@vm18042test:/home/user# efa system supportsave
SupportSave File Location: /var/log/efa/efa_2025-05-24T09-03-21.378.logs.zip
SupportSave File Location in Remote Server: 10.32.85.5:/home/user/Downloads/
efa_2024-09-24T09-03-21.378.logs.zip
```

- **Backup**

  To back up the system and copy the backup to a remote server (after configuring remote server settings):

```
(efa:user)root@vm18042test:/home/user# efa system backup --remote
Generating backup of EFA...
Backup Location on local server: /var/log/efa/backup/
EFA-3.8.0-1109-2025-03-24T09-05-39.002.tar
Backup Location on Remote Server: root@10.32.85.5:/home/user/Downloads/
EFA-3.8.0-1109-2025-03-24T09-05-39.002.tar
```

Enable or Disable Flooding for IP DHCP Relay

You can enable or disable flooding for IP DHCP relay.

About This Task
Follow this procedure to enable or disable IP DHCP relay flooding on a device.

Note
Ensure that you are using SLX firmware version 20.5.3 or later.

Procedure

Run the following command to enable or disable IP DHCP relay flooding on a device:

```
efa fabric show

Fabric Name: default, Fabric Description: Default Fabric, 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 |
+------------+-----+-----------+-----+------+--------------+-----------+------------------
+-----------------+---------+-------+
+------------+-----+-----------+-----+------+--------------+-----------+------------------
+-----------------+---------+-------+

Fabric Name: fs, Fabric Description: , Fabric Type: non-clos, Fabric
Status: configure-success, Fabric Health: Green
+--------------+-----+---------------+------------+------+--------------+-------------
+------------------+-----------------+---------+-------+
| IP ADDRESS   |RACK | HOST NAME     |    ASN
| ROLE | DEVICE STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-----+---------------+------------+------+--------------+-------------
+------------------+-----------------+---------+-------+
| 10.20.246.15 | r1  | Rack1-Device1 | 4200000000
| Leaf | provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
| 10.20.246.16 | r1  | Rack1-Device2 | 4200000000
| Leaf | provisioned  | cfg in-sync | NA                | NA              | 2       | 1     |
+--------------+-----+---------------+------------+------+--------------+-------------
+------------------+-----------------+---------+-------+

efa inventory device setting update --ip 10.20.246.15 --ip-dhcp-relay-disable-flooding No
+-------------+---------------------------------+---------+-------+-------+
| IP ADDRESS  |              NAME               | STATUS  | VALUE | ERROR |
+-------------+---------------------------------+---------+-------+-------+
| 10.20.246.15| Ip Dhcp Relay Disable Flooding | Success | No    |       |
+-------------+---------------------------------+---------+-------+-------+

efa inventory device setting update --ip-dhcp-relay-disable-flooding No --ip 10.20.246.15,10.20.246.16
+-------------+---------------------------------+---------+-------+-------+
| IP ADDRESS  |              NAME               | STATUS  | VALUE | ERROR |
+-------------+---------------------------------+---------+-------+-------+
| 10.20.246.15| Ip Dhcp Relay Disable Flooding | Success | No    |       |
+-------------+---------------------------------+---------+-------+-------+
| 10.20.246.16| Ip Dhcp Relay Disable Flooding | Success | No    |       |
+-------------+---------------------------------+---------+-------+-------+


efa inventory device setting update --ip-dhcp-relay-disable-flooding Yes --fabric fs
+-------------+---------------------------------+---------+-------+-------+
| IP ADDRESS  |              NAME               | STATUS  | VALUE | ERROR |
+-------------+---------------------------------+---------+-------+-------+
| 10.20.246.15| Ip Dhcp Relay Disable Flooding | Success | Yes   |       |
+-------------+---------------------------------+---------+-------+-------+
| 10.20.246.16| Ip Dhcp Relay Disable Flooding | Success | Yes   |       |
+-------------+---------------------------------+---------+-------+-------+


efa inventory device setting show [flags]
```

```
Flags:
  --ip string      IP address of device
```

```
efa inventory device setting show --ip 10.20.246.15           efa
inventory device setting show --ip 10.20.246.16
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
|              NAME              | VALUE |  APP STATE  |
|              NAME              | VALUE |  APP STATE  |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Maintenance Mode Enable On     | No    |             |       |
Maintenance Mode Enable On     | No    |             |
| Reboot                         |       |             |       |
Reboot                         |       |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Maintenance Mode Enable        | No    |             |       |
Maintenance Mode Enable        | No    |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Maintenance Convergence Time   |       |             |       |
Maintenance Convergence Time   |       |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| MCT Bring-up Delay             |       |             |       | MCT
Bring-up Delay             |       |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Health Check Enabled           | No    |             |       | Health
Check Enabled           | No    |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Health Check Interval          | 6m    |             |       | Health
Check Interval          | 6m    |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Health Check Heartbeat Miss    | 2     |             |       | Health
Check Heartbeat Miss    | 2     |             |
| Threshold                      |       |             |       |
Threshold                      |       |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Periodic Backup Enabled        | Yes   |             |       |
Periodic Backup Enabled        | Yes   |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Config Backup Interval         | 24h   |             |       | Config
Backup Interval         | 24h   |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Config Backup Count            | 4     |             |       | Config
Backup Count            | 4     |             |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Prefix Independent Convergence | No    | cfg-in-sync |       | Prefix
Independent Convergence | No    | cfg-in-sync |
+--------------------------------+-------+-------------+
+--------------------------------+-------+-------------+
| Static Prefix Independent      | No    |             |       | Static
Prefix Independent      | No    |             |
| Convergence                    |       |             |       |
```

```
Convergence                    |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Maximum Load Sharing Paths   |       |             |      |
Maximum Load Sharing Paths   |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Maximum Ipv6 Prefix Length 64  |       |             |      |
Maximum Ipv6 Prefix Length 64  |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Urpf                         |       |             |      |
Urpf                         |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Ip Dhcp Relay Disable Flooding | Yes   | cfg-in-sync |      | Ip
Dhcp Relay Disable Flooding | Yes   | cfg-in-sync |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Ip Option Disable              | Yes   | cfg-in-sync |      | Ip
Option Disable              | Yes   | cfg-in-sync |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Ip Option Disable Cpu          | No    | cfg-in-sync |      | Ip
Option Disable Cpu          | No    | cfg-in-sync |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Ipv6 Option Disable            | No    | cfg-in-sync |      | Ipv6
Option Disable            | No    | cfg-in-sync |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Crypto Certificate Expiry Info |       |             |      | Crypto
Certificate Expiry Info |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Crypto Certificate Expiry      |       |             |      | Crypto
Certificate Expiry      |       |             |
| Minor                        |       |             |      |
Minor                        |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Crypto Certificate Expiry      |       |             |      | Crypto
Certificate Expiry      |       |             |
| Major                        |       |             |      |
Major                        |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Crypto Certificate Expiry      |       |             |      | Crypto
Certificate Expiry      |       |             |
| Critical                     |       |             |      |
Critical                     |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
| Peer Group Ipv6 Prefix Over    |       |             |      | Peer
Group Ipv6 Prefix Over    |       |             |
| Ipv4 Peer                    |       |             |      | Ipv4
Peer                    |       |             |
+------------------------------+-------+-------------+
+------------------------------+-------+-------------+
```

```
Rack1-Device1# show running ip dhcp
ip dhcp relay disable-flooding
```

```
Rack1-Device2# show running ip dhcp
ip dhcp relay disable-flooding
```

Show Device Adapter Connection Status

You can view the connection status of device adapters.

About This Task
Follow this procedure to view the connection status of device adapters.

Device adapters are utilized to connect to registered devices. Monitoring the status of these adapters is helpful in identifying connectivity issues.

Procedure

To show the connection status of device adapters, run the following command:

```
efa inventory debug device-adapter-status
```

Note
For information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Example

```
efa inventory debug device-adapter-status --ip 10.24.80.58
+-------------+-------------------+----------+-----------+--------------+
|     IP      |      Adapter      | Protocol |   State   | Error Status |
+-------------+-------------------+----------+-----------+--------------+
| 10.24.80.58 | SwitchRestAdapter |  https   | Connected |   Success    |
+-------------+-------------------+----------+-----------+--------------+
| 10.24.80.58 | SwitchSSHAdapter  |   SSH    | Connected |   Success    |
+-------------+-------------------+----------+-----------+--------------+
| 10.24.80.58 |  SwitchNcAdapter  | NetConf  | Connected |   Success    |
+-------------+-------------------+----------+-----------+--------------+
Device Adapter Status
--- Time Elapsed: 5.867081913s ---
```

Show Device Certificate Expiry Time

You can view the status of device certificates and their expiry time.

About This Task
Follow this procedure to check the status and expiry time of device certificates.

During the registration process, devices are equipped with certificates that are used for connectivity logging. These certificates hold important information such as their type, expiration date, and status, including whether they are synchronized with XCO certificates. These details are particularly useful when debugging connectivity issues.

Procedure

To show the device certificate status including expiration time, run the following command:

```
efa certificate device expiry show
```

Note
For information about commands and supported parameters to see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Example

```
efa certificate device expiry show  --ip 10.139.44.147
+---------------+----------+-------------------------------+------------------------
+----------+
|      IP       |   Type   |            Expiry             | EFA/Device
Sync Status | Status  |
+---------------+----------+-------------------------------+------------------------
+----------+
| 10.139.44.147 |  HTTPS   | 2025-08-06 18:11:47 +0000 GMT |
In Sync         | Success |
+---------------+----------+-------------------------------+------------------------
+----------+
| 10.139.44.147 | SyslogCA | 2033-09-03 17:58:17 +0000 GMT |
In Sync         | Success |
+---------------+----------+-------------------------------+------------------------
+----------+
| 10.139.44.147 |  OAuth2  | 2033-09-03 18:07:50 +0000 GMT |
In Sync         | Success |
+---------------+----------+-------------------------------+------------------------
+----------+
Device Certificate Expiry Show
--- Time Elapsed: 53.13444312s ---
```

Configure Device Certificate Expiry Time

You can configure alerts for a TLS certificate expiration.

About This Task

Follow this procedure to configure alerts for a TLS certificate expiration in SLX.

The configuration enables SLX to send the alerts via RASLOG and SNMP traps. You can configure the following severity level along with the time period:

- Critical
- Major
- Minor
- Info

The time period for generating SNMP traps corresponds to the input time in days before certificate expiry. The input range is from 1 to 90 days, with a 0 value used to de-configure.

Procedure

To configure the device certificates expiry time, run the following command:

```
efa inventory device setting update --crypto-cert-expiry-info
efa inventory device setting update --crypto-cert-expiry-minor
efa inventory device setting update --crypto-cert-expiry-major
efa inventory device setting update --crypto-cert-expiry-critical
```

Note

For information about commands and supported parameters to see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Table 22: Drift Reconcile and Idempotency Support

Identify Drift	Reconcile Configuration	Idempotency
Yes	Yes	Yes

Example

```
efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-
info 50
+---------------+---------------------------------+---------+-------+-------+
|  IP ADDRESS   |              NAME               | STATUS  | VALUE | ERROR |
+---------------+---------------------------------+---------+-------+-------+
| 10.20.24.10   | Crypto Cert Expiry Info         | Success |  50   |       |
+---------------+---------------------------------+---------+-------+-------+

efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-
minor 30
+---------------+---------------------------------+---------+-------+-------+
|  IP ADDRESS   |              NAME               | STATUS  | VALUE | ERROR |
+---------------+---------------------------------+---------+-------+-------+
| 10.20.24.10   | Crypto Cert Expiry Minor        | Success |  30   |       |
+---------------+---------------------------------+---------+-------+-------+

efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-
major 10
+---------------+---------------------------------+---------+-------+-------+
|  IP ADDRESS   |              NAME               | STATUS  | VALUE | ERROR |
+---------------+---------------------------------+---------+-------+-------+
| 10.20.24.10   | Crypto Cert Expiry Major        | Success |  10   |       |
+---------------+---------------------------------+---------+-------+-------+

efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-
critical 5
+---------------+---------------------------------+---------+-------+-------+
|  IP ADDRESS   |              NAME               | STATUS  | VALUE | ERROR |
+---------------+---------------------------------+---------+-------+-------+
| 10.20.24.10   | Crypto Cert Expiry Critical     | Success |   5   |       |
+---------------+---------------------------------+---------+-------+-------+
```

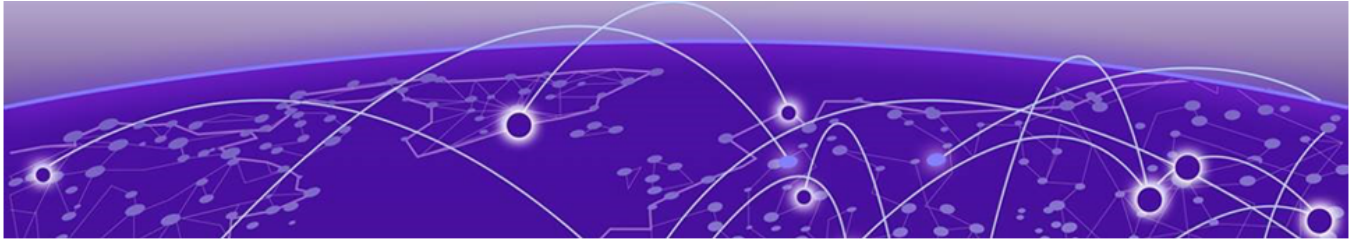

XCO Event Management

Learn about RASlog, SNMP, and Notification services.

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, 4.1.0*.

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.
- As a best practice, use the rsyslog server version 8.36.0 and above. Avoid the following issue in 8.36.0:

  ```
  Oversized messages causes server close/reset connection.
  ```

  In the syslog version lower than 8.36.0, such notifications are dropped.

Notification Types

XCO notifies the subscribers of 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 have more alert-specific data.

- APP_ALARMS – Alarms are generated by XCO.
- HOST_EVENTS – Events generated by the HOST OS

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.
- The HOST_EVENTS filter and its sub filter host-event are mandatory fields for auditd. The following is the minimum required CLI command to capture host events:

```
efa notification
subscribers add-syslog-relp --address 10.37.11.38:20514 --insecure --filter
HOST_EVENTS --host-event auditd
```

Default Behavior: If no key and record type are specified, XCO will default to sending only USER_LOGIN type audit logs to the external server as shown in the following example:

```
efa notification subscribers add-syslog-relp
--address 10.37.11.38:20514 --insecure --filter HOST_EVENTS --host-event
auditd
```

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.

Abbreviation	Expansion
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.
Security	Audit all user-initiated security events across 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, 4.1.0*

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 3.3.0) 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 3.3.0, changes to the external webhook or syslog server are not necessary.

XCO Notification Tasks

XCO task represents an XCO Notification task with an ID and message pattern type XCO Task struct {"ID", "string"}.

Use this topic to learn about the following service notification tasks:

- Fabric Service Notifications on page 831
- Tenant Service Notifications on page 833
- Asset Service Notifications on page 838

- Policy Service Notifications on page 850
- Notifications Service on page 854

Fabric Service Notifications

Note
%S is fabric name.

XCO Notification Type	XCO Task
FabricCreateRequestRCVD	{"EFA-000001", "Fabric create request received :request=%s"}
FabricCreateRequestSuccess	{"EFA-000002", "Fabric create request success :request=%s"}
FabricCreateRequestFailed	{"EFA-000003", "Fabric create request failed :request=%s,Error=%s"}
FabricDeleteRequestRCVD	{"EFA-000004", "Fabric delete request received :request=%s"}
FabricDeleteRequestSuccess	{"EFA-000005", "Fabric delete request success :request=%s"}
FabricDeleteRequestFailed	{"EFA-000006", "Fabric delete request failed :request=%s,Error=%s"}
FabricCloneRequestRCVD	{"EFA-000007", "Fabric clone request received :request=%s"}
FabricCloneRequestSuccess	{"EFA-000008", "Fabric clone request success :request=%s"}
FabricCloneRequestFailed	{"EFA-000009", "Fabric clone request failed :request=%s,Error=%s"}
FabricDeviceAddRequestRCVD	{"EFA-000010", "Fabric device add request received :request=%s"}
FabricDeviceAddRequestSuccess	{"EFA-000011", "Fabric device add request success :request=%s"}
FabricDeviceAddRequestFailed	{"EFA-000012", "Fabric device add request failed :request=%s,Error=%s"}
FabricDeviceDeleteRequestRCVD	{"EFA-000013", "Fabric device delete request received :request=%s"}
FabricDeviceDeleteRequestSuccess	{"EFA-000014", "Fabric device delete request success :request=%s"}
FabricDeviceDeleteRequestFailed	{"EFA-000015", "Fabric device delete request failed :request=%s,Error=%s"}
FabricConfigureRequestRCVD	{"EFA-000016", "Fabric configure request received :request=%s"}

XCO Notification Type	XCO Task
FabricConfigureRequestSuccess	{"EFA-000017", "Fabric configure request success :request=%s"}
FabricConfigureRequestFailed	{"EFA-000018", "Fabric configure request failed :request=%s,Error=%s"}
FabricImportRequestRCVD	{"EFA-000019", "Fabric import request received :request=%s"}
FabricImportRequestSuccess	{"EFA-000020", "Fabric import request success :request=%s"}
FabricImportRequestFailed	{"EFA-000021", "Fabric import request failed :request=%s,Error=%s"}
FabricDRRequestRCVD	{"EFA-000022", "Fabric drift-reconcile request received :request=%s"}
FabricDRRequestSuccess	{"EFA-000023", "Fabric drift-reconcile request success :request=%s"}
FabricDRRequestFailed	{"EFA-000024", "Fabric drift-reconcile request failed :request=%s,Error=%s"}
FabricSettingUpdateRequestRCVD	{"EFA-000025", "Fabric settings update request received :request=%s"}
FabricSettingUpdateRequestSuccess	{"EFA-000026", "Fabric settings update request success :request=%s"}
FabricSettingUpdateRequestFailed	{"EFA-000027", "Fabric settings update request failed :request=%s,Error=%s"}
FabricClearConfigRequestRCVD	{"EFA-000028", "Fabric clear config request received :request=%s"}
FabricClearConfigRequestSuccess	{"EFA-000029", "Fabric clear config request success :request=%s"}
FabricClearConfigRequestFailed	{"EFA-000030", "Fabric clear config request failed :request=%s,Error=%s"}
FabricDeleteExecutionsRequestRCVD	{"EFA-000031", "Fabric delete executions request received :request=%s"}
FabricDeleteExecutionsRequestSuccess	{"EFA-000032", "Fabric delete executions request success :request=%s"}
FabricDeleteExecutionsRequestFailed	{"EFA-000033", "Fabric delete executions request failed :request=%s,Error=%s"}
FabricValidationRequestRCVD	{"EFA-000034", "Fabric validation executions request received :request=%s"}
FabricValidationRequestSuccess	{"EFA-000035", "Fabric validation executions request success :request=%s"}
FabricValidationRequestFailed	{"EFA-000036", "Fabric validation executions request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
FabricShowRequestRCVD	{"EFA-000037", "Fabric show executions request received :request=%s"}
FabricShowRequestSuccess	{"EFA-000038", "Fabric show executions request success :request=%s"}
FabricShowRequestFailed	{"EFA-000039", "Fabric show executions request failed :request=%s,Error=%s"}
FabricMigrateRequestRCVD	{"EFA-000040", "Fabric migrate request received :request=%s"}
FabricMigrateRequestSuccess	{"EFA-000041", "Fabric migrate request success :request=%s"}
FabricMigrateRequestFailed	{"EFA-000042", "Fabric migrate request failed :request=%s,Error=%s"}

Tenant Service Notifications

Note
%S is tenant name.

XCO Notification Type	XCO Task
TenantCreateRequestRCVD	{"EFA-001001", "Tenant create request received :request=%s"}
TenantCreateRequestSuccess	{"EFA-001002", "Tenant create request success :request=%s"}
TenantCreateRequestFailed	{"EFA-001003", "Tenant create request failed :request=%s,Error=%s"}
TenantUpdateRequestRCVD	{"EFA-001004", "Tenant update request received :request=%s"}
TenantUpdateRequestSuccess	{"EFA-001005", "Tenant update request success :request=%s"}
TenantUpdateRequestFailed	{"EFA-001006", "Tenant update request failed :request=%s,Error=%s"}
TenantDeleteRequestRCVD	{"EFA-001007", "Tenant delete request received :request=%s"}
TenantDeleteRequestSuccess	{"EFA-001008", "Tenant delete request success :request=%s"}
TenantDeleteRequestFailed	{"EFA-001009", "Tenant delete request failed :request=%s,Error=%s"}
TenantVRFCreateRequestRCVD	{"EFA-001010", "Tenant vrf create request received :request=%s"}
TenantVRFCreateRequestSuccess	{"EFA-001011", "Tenant vrf create request success :request=%s"}

XCO Notification Type	XCO Task
TenantVRFCreateRequestFailed	{"EFA-001012", "Tenant vrf create request failed :request=%s,Error=%s"}
TenantVRFUpdateRequestRCVD	{"EFA-001013", "Tenant vrf update request received :request=%s"}
TenantVRFUpdateRequestSuccess	{"EFA-001014", "Tenant vrf update request success :request=%s"}
TenantVRFUpdateRequestFailed	{"EFA-001015", "Tenant vrf update request failed :request=%s,Error=%s"}
TenantVRFDeleteRequestRCVD	{"EFA-001016", "Tenant vrf delete request received :request=%s"}
TenantVRFDeleteRequestSuccess	EFA-001017", "Tenant vrf delete request success :request=%s"}
TenantVRFDeleteRequestFailed	{"EFA-001018", "Tenant vrf delete request failed :request=%s,Error=%s"}
TenantVRFConfigureRequestRCVD	{"EFA-001013", "Tenant vrf configure request received :request=%s"}
TenantVRFConfigureRequestSuccess	{"EFA-001014", "Tenant vrf configure request success :request=%s"}
TenantVRFConfigureRequestFailed	{"EFA-001015", "Tenant vrf configure request failed :request=%s,Error=%s"}
TenantEPGCreateRequestRCVD	{"EFA-001019", "Tenant epg create request received :request=%s"}
TenantEPGCreateRequestSuccess	{"EFA-001020", "Tenant epg create request success :request=%s"}
TenantEPGCreateRequestFailed	{"EFA-001021", "Tenant epg create request failed :request=%s,Error=%s"}
TenantEPGUpdateRequestRCVD	{"EFA-001022", "Tenant epg update request received :request=%s"}
TenantEPGUpdateRequestSuccess	{"EFA-001023", "Tenant epg update request success :request=%s"}
TenantEPGUpdateRequestFailed	{"EFA-001024", "Tenant epg update request failed :request=%s,Error=%s"}
TenantEPGDeleteRequestRCVD	{"EFA-001025", "Tenant epg delete request received :request=%s"}
TenantEPGDeleteRequestSuccess	{"EFA-001026", "Tenant epg delete request success :request=%s"}
TenantEPGDeleteRequestFailed	{"EFA-001027", "Tenant epg delete request failed :request=%s,Error=%s"}
TenantEPGConfigureRequestRCVD	{"EFA-001028", "Tenant epg configure request received :request=%s"}

XCO Notification Type	XCO Task
TenantEPGConfigureRequestSuccess	{"EFA-001029", "Tenant epg configure request success :request=%s"}
TenantEPGConfigureRequestFailed	{"EFA-001030", "Tenant epg configure request failed :request=%s,Error=%s"}
TenantEPGDetachRequestRCVD	{"EFA-001031", "Tenant epg detach request received :request=%s"}
TenantEPGDetachRequestSuccess	{"EFA-001032", "Tenant epg detach request success :request=%s"}
TenantEPGDetachRequestFailed	{"EFA-001033", "Tenant epg detach request failed :request=%s,Error=%s"}
TenantPOCreateRequestRCVD	{"EFA-001034", "Tenant po create request received :request=%s"}
TenantPOCreateRequestSuccess	{"EFA-001035", "Tenant po create request success :request=%s"}
TenantPOCreateRequestFailed	{"EFA-001036", "Tenant po create request failed :request=%s,Error=%s"}
TenantPOUpdateRequestRCVD	{"EFA-001037", "Tenant po update request received :request=%s"}
TenantPOUpdateRequestSuccess	{"EFA-001038", "Tenant po update request success :request=%s"}
TenantPOUpdateRequestFailed	{"EFA-001039", "Tenant po update request failed :request=%s,Error=%s"}
TenantPODeleteRequestRCVD	{"EFA-001040", "Tenant po delete request received :request=%s"}
TenantPODeleteRequestSuccess	{"EFA-001041", "Tenant po delete request success :request=%s"}
TenantPODeleteRequestFailed	{"EFA-001042", "Tenant po delete request failed :request=%s,Error=%s"}
TenantPOConfigureRequestRCVD	{"EFA-001043", "Tenant po configure request received :request=%s"}
TenantPOConfigureRequestSuccess	{"EFA-001044", "Tenant po configure request success :request=%s"}
TenantPOConfigureRequestFailed	{"EFA-001045", "Tenant po configure request failed :request=%s,Error=%s"}
TenantDeleteExecutionsRequestRCVD	{"EFA-001046", "Tenant delete executions request received :request=%s"}
TenantDeleteExecutionsRequestSuccess	{"EFA-001047", "Tenant delete executions request success :request=%s"}
TenantDeleteExecutionsRequestFailed	{"EFA-001048", "Tenant delete executions request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
TenantDRRequestRCVD	{"EFA-001049", "Tenant drift-reconcile request received :request=%s"}
TenantDRRequestSuccess	{"EFA-001050", "Tenant drift-reconcile request success :request=%s"}
TenantDRRequestFailed	{"EFA-001051", "Tenant drift-reconcile request failed :request=%s,Error=%s"}
TenantBGPCreateRequestRCVD	{"EFA-001052", "Tenant bgp create request received :request=%s"}
TenantBGPCreateRequestSuccess	{"EFA-001053", "Tenant bgp create request success :request=%s"}
TenantBGPCreateRequestFailed	{"EFA-001054", "Tenant bgp create request failed :request=%s,Error=%s"}
TenantBGPUpdateRequestRCVD	{"EFA-001055", "Tenant bgp update request received :request=%s"}
TenantBGPUpdateRequestSuccess	{"EFA-001056", "Tenant bgp update request success :request=%s"}
TenantBGPUpdateRequestFailed	{"EFA-001057", "Tenant bgp update request failed :request=%s,Error=%s"}
TenantBGPDeleteRequestRCVD	{"EFA-001058", "Tenant bgp delete request received :request=%s"}
TenantBGPDeleteRequestSuccess	{"EFA-001059", "Tenant bgp delete request success :request=%s"}
TenantBGPDeleteRequestFailed	{"EFA-001060", "Tenant bgp delete request failed :request=%s,Error=%s"}
TenantBGPConfigureRequestRCVD	{"EFA-001061", "Tenant bgp configure request received :request=%s"}
TenantBGPConfigureRequestSuccess	{"EFA-001062", "Tenant bgp configure request success :request=%s"}
TenantBGPConfigureRequestFailed	{"EFA-001063", "Tenant bgp configure request failed :request=%s,Error=%s"}
TenantBGPServicePeerGroupCreateRequestRCVD	{"EFA-001064", "Tenant bgp service peer-group create request received :request=%s"}
TenantBGPServicePeerGroupCreateRequestSuccess	{"EFA-001065", "Tenant bgp service peer-group create request success :request=%s"}
TenantBGPServicePeerGroupCreateRequestFailed	{"EFA-001066", "Tenant bgp service peer-group create request failed :request=%s,Error=%s"}
TenantBGPServicePeerGroupUpdateRequestRCVD	{"EFA-001067", "Tenant bgp service peer-group update request received :request=%s"}
TenantBGPServicePeerGroupUpdateRequestSuccess	{"EFA-001068", "Tenant bgp service peer-group update request success :request=%s"}

XCO Notification Type	XCO Task
TenantBGPServicePeerGroupUpdateRequestFailed	{"EFA-001069", "Tenant bgp service peer-group update request failed :request=%s,Error=%s"}
TenantBGPServicePeerGroupDeleteRequestRCVD	{"EFA-001070", "Tenant bgp service peer-group delete request received :request=%s"}
TenantBGPServicePeerGroupDeleteRequestSuccess	{"EFA-001071", "Tenant bgp service peer-group delete request success :request=%s"}
TenantBGPServicePeerGroupDeleteRequestFailed	{"EFA-001072", "Tenant bgp service peer-group delete request failed :request=%s,Error=%s"}
TenantBGPServicePeerGroupConfigureRequestRCVD	{"EFA-001073", "Tenant bgp service peer-group configure request received :request=%s"}
TenantBGPServicePeerGroupConfigureRequestSuccess	{"EFA-001074", "Tenant bgp service peer-group configure request success :request=%s"}
TenantBGPServicePeerGroupConfigureRequestFailed	{"EFA-001075", "Tenant bgp service peer-group configure request failed :request=%s,Error=%s"}
TenantMirrorServiceSessionCreateRequestRCVD	{"EFA-001076", "Tenant mirror service session create request received :request=%s"}
TenantMirrorServiceSessionCreateRequestSuccess	{"EFA-001077", "Tenant mirror service session create request success :request=%s"}
TenantMirrorServiceSessionCreateRequestFailed	{"EFA-001078", "Tenant mirror service session create request failed :request=%s,Error=%s"}
TenantMirrorServiceSessionDeleteRequestRCVD	{"EFA-001079", "Tenant mirror service session delete request received :request=%s"}
TenantMirrorServiceSessionDeleteRequestSuccess	{"EFA-001080", "Tenant mirror service session delete request success :request=%s"}
TenantMirrorServiceSessionDeleteRequestFailed	{"EFA-001081", "Tenant mirror service session delete request failed :request=%s,Error=%s"}
TenantMirrorServiceSessionConfigureRequestRCVD	{"EFA-001082", "Tenant mirror service session configure request received :request=%s"}
TenantMirrorServiceSessionConfigureRequestSuccess	{"EFA-001083", "Tenant mirror service session configure request success :request=%s"}
TenantMirrorServiceSessionConfigureRequestFailed	{"EFA-001084", "Tenant mirror service session configure request failed :request=%s,Error=%s"}

Asset Service Notifications

Note
%S is an Asset name.

XCO Notification Type	XCO Task
AssetDeviceDeleteRequestRCVD	{"EFA-002001", "Asset device delete request received :request=%s"}
AssetDeviceDeleteRequestSuccess	{"EFA-002002", "Asset device delete request success :request=%s"}
AssetDeviceDeleteRequestFailed	{"EFA-002003", "Asset device delete request failed :request=%s,Error=%s"}
AssetRegisterSwitchRequestRCVD	{"EFA-002004", "Asset register switch request received :request=%s"}
AssetRegisterSwitchRequestSuccess	{"EFA-002005", "Asset register switch request success :request=%s"}
AssetRegisterSwitchRequestFailed	{"EFA-002006", "Asset register switch request failed :request=%s,Error=%s"}
AssetUpdateSwitchRequestRCVD	{"EFA-002007", "Asset update switch request received :request=%s"}
AssetUpdateSwitchRequestSuccess	{"EFA-002008", "Asset update switch request success :request=%s"}
AssetUpdateSwitchRequestFailed	{"EFA-002009", "Asset update switch request failed :request=%s,Error=%s"}
AssetReplaceSwitchRequestRCVD	{"EFA-002010", "Asset replace switch request received :request=%s"}
AssetReplaceSwitchRequestSuccess	{"EFA-002011", "Asset replace switch request success :request=%s"}
AssetReplaceSwitchRequestFailed	{"EFA-002012", "Asset replace switch request failed :request=%s,Error=%s"}
AssetCompareSwitchRequestRCVD	{"EFA-002013", "Asset compare switch request received :request=%s"}
AssetCompareSwitchRequestSuccess	{"EFA-002014", "Asset compare switch request success :request=%s"}
AssetCompareSwitchRequestFailed	{"EFA-002015", "Asset compare switch request failed :request=%s,Error=%s"}
AssetManageSwitchesConfigRequestRCVD	{"EFA-002016", "Asset manage switches config request received :request=%s"}
AssetManageSwitchesConfigRequestSuccess	{"EFA-002017", "Asset manage switches config request success :request=%s"}
AssetManageSwitchesConfigRequestFailed	{"EFA-002018", "Asset manage switches config request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
AssetExecuteSwitchesCliRequestRCVD	{"EFA-002019", "Asset execute switches CLI request received :request=%s"}
AssetExecuteSwitchesCliRequestSuccess	{"EFA-002020", "Asset execute switches CLI request success :request=%s"}
AssetExecuteSwitchesCliRequestFailed	{"EFA-002021", "Asset execute switches CLI request failed :request=%s,Error=%s"}
AssetStartConfigBackupRequestRCVD	{"EFA-002022", "Asset create config backup request received :request=%s"}
AssetStartConfigBackupRequestSuccess	{"EFA-002023", "Asset create config backup request success :request=%s"}
AssetStartConfigBackupRequestFailed {	"EFA-002024", "Asset create config backup failed :request=%s,Error=%s"}
AssetDeleteConfigBackupRequestRCVD	{"EFA-002025", "Asset delete config backup request received :request=%s"}
AssetStartConfigReplayRequestSuccess	{"EFA-002026", "Asset delete config backup request success :request=%s"}
AssetDeleteConfigBackupRequestFailed	{"EFA-002027", "Asset delete config backup request failed :request=%s,Error=%s"}
AssetStartConfigReplayRequestRCVD	{"EFA-002028", "Asset start config replay request received :request=%s"}
AssetStartConfigReplayRequestSuccess	{"EFA-002029", "Asset start config replay request success :request=%s"}
AssetStartConfigReplayRequestFailed	{"EFA-002030", "Asset start config replay request failed :request=%s,Error=%s"}
AssetDeleteConfigReplayRequestRCVD	{"EFA-002031", "Asset delete config replay request received :request=%s"}
AssetDeleteConfigReplayRequestSuccess	{"EFA-002032", "Asset delete config replay request success :request=%s"}
AssetDeleteConfigReplayRequestFailed	{"EFA-002033", "Asset delete config replay request failed :request=%s,Error=%s"}
AssetCreateBreakoutRequestRCVD	{"EFA-002034", "Asset create breakout request received :request=%s"}
AssetCreateBreakoutRequestSuccess	{"EFA-002035", "Asset create breakout request success :request=%s"}
AssetSetInterfaceMtuRequestFailed	{"EFA-002036", "Asset create breakout request failed :request=%s,Error=%s"}
AssetRemoveBreakoutRequestRCVD	{"EFA-002037", "Asset remove breakout request received :request=%s"}
AssetRemoveBreakoutRequestSuccess	{"EFA-002038", "Asset remove breakout request success :request=%s"}

XCO Notification Type	XCO Task
AssetRemoveBreakoutRequestFailed	{"EFA-002039", "Asset remove breakout request failed :request=%s,Error=%s"}
AssetRemoveBreakoutRequestFailed	{"EFA-002040", "Asset change admin state request received :request=%s"}
AssetChangeAdminStateRequestSuccess	{"EFA-002041", "Asset change admin state request success :request=%s"}
AssetChangeAdminStateRequestFailed	{"EFA-002042", "Asset change admin state request failed :request=%s,Error=%s"}
AssetChangePortSpeedRequestRCVD	{"EFA-002043", "Asset change port speed request received :request=%s"}
AssetChangePortSpeedRequestSuccess	{"EFA-002044", "Asset change port speed request success :request=%s"}
AssetChangePortSpeedRequestFailed	{"EFA-002045", "Asset change port speed request failed :request=%s,Error=%s"}
AssetSetInterfaceMtuRequestRCVD	{"EFA-002046", "Asset set interface mtu request received :request=%s"}
AssetChangePortSpeedRequestFailed	{"EFA-002047", "Asset set interface mtu request success :request=%s"}
AssetSetInterfaceMtuRequestFailed	{"EFA-002048", "Asset set interface mtu request failed :request=%s,Error=%s"}
AssetUpdateKeyStoreRequestRCVD	{"EFA-002049", "Asset update key store request received :request=%s"}
AssetUpdateKeyStoreRequestSuccess	{"EFA-002050", "Asset update key store request success :request=%s"}
AssetUpdateKeyStoreRequestFailed	{"EFA-002051", "Asset update key store request failed :request=%s,Error=%s"}
AssetDeleteKeyStoreRequestRCVD	{"EFA-002052", "Asset delete key store request received :request=%s"}
AssetDeleteKeyStoreRequestSuccess	{"EFA-002053", "Asset delete key store request success :request=%s"}
AssetDeleteKeyStoreRequestFailed {	"EFA-002054", "Asset delete key store request failed :request=%s,Error=%s"}
AssetRegisterFirmwareHostRequestRCVD	{"EFA-002055", "Asset register firmware host request received :request=%s"}
AssetRegisterFirmwareHostRequestSuccess	{"EFA-002056", "Asset register firmware host request success :request=%s"}
AssetRegisterFirmwareHostRequestFailed	{"EFA-002057", "Asset register firmware host request failed :request=%s,Error=%s"}
AssetDeleteFirmwareHostRequestRCVD	{"EFA-002058", "Asset delete firmware host request received :request=%s"}

XCO Notification Type	XCO Task
AssetDeleteFirmwareHostRequestSuccess	{"EFA-002059", "Asset delete firmware host request success :request=%s"}
AssetDeleteFirmwareHostRequestFailed	{"EFA-002060", "Asset delete firmware host request failed :request=%s,Error=%s"}
AssetUpdateFirmwareHostRequestRCVD	{"EFA-002061", "Asset update firmware host request received :request=%s"}
AssetUpdateFirmwareHostRequestSuccess	{"EFA-002062", "Asset update firmware host request success :request=%s"}
AssetUpdateFirmwareHostRequestFailed	{"EFA-002063", "Asset update firmware host request failed :request=%s,Error=%s"}
AssetAddPreparedSwitchesRequestRCVD	{"EFA-002064", "Asset add prepared switches request received :request=%s"}
AssetAddPreparedSwitchesRequestSuccess	{"EFA-002065", "Asset add prepared switches request success :request=%s"}
AssetAddPreparedSwitchesRequestFailed	{"EFA-002066", "Asset add prepared switches request failed :request=%s,Error=%s"}
AssetRemovePreparedSwitchesRequestRCVD	{"EFA-002067", "Asset remove prepared switches request received :request=%s"}
AssetUpdateDeviceDetailsRequestFailed	{"EFA-002068", "Asset remove prepared switches request success :request=%s"}
AssetRemovePreparedSwitchesRequestFailed	{"EFA-002069", "Asset remove prepared switches request failed :request=%s,Error=%s"}
AssetExecuteFirmwareDownloadRequestRCVD	{"EFA-002070", "Asset execute firmware download request received :request=%s"}
AssetExecuteFirmwareDownloadRequestSuccess	{"EFA-002071", "Asset execute firmware download request success :request=%s"}
AssetExecuteFirmwareDownloadRequestFailed	{"EFA-002072", "Asset execute firmware download request failed :request=%s,Error=%s"}
AssetUpdateSwitchSettingRequestRCVD	{"EFA-002073", "Asset update switch setting request received :request=%s"}
AssetUpdateSwitchSettingRequestSuccess	{"EFA-002074", "Asset update switch setting request success :request=%s"}
AssetUpdateSwitchSettingRequestFailed	{"EFA-002075", "Asset update switch setting request failed :request=%s,Error=%s"}
AssetDriftAndReconcileExecuteHandlerRCVD	{"EFA-002076", "Asset drift-reconcile execute request received :request=%s"}
AssetDriftAndReconcileExecuteHandlerSuccess	{"EFA-002077", "Asset drift-reconcile execute request success :request=%s"}

XCO Notification Type	XCO Task
AssetDriftAndReconcileExecuteHandlerFailed	{"EFA-002078", "Asset drift-reconcile execute request failed :request=%s,Error=%s"}
AssetDeleteDriftAndReconcileRequestRCVD	{"EFA-002079", "Asset drift-reconcile request received :request=%s"}
AssetDeleteDriftAndReconcileRequestSuccess	{"EFA-002080", "Asset drift-reconcile request success :request=%s"}
AssetDeleteDriftAndReconcileRequestFailed	{"EFA-002081", "Asset drift-reconcile request failed :request=%s,Error=%s"}
AssetDeleteRMARequestRCVD	{"EFA-002082", "Asset delete rma request received :request=%s"}
AssetDeleteRMARequestSuccess	{"EFA-002083", "Asset delete rma request success :request=%s"}
AssetDeleteRMARequestFailed	{"EFA-002084", "Asset delete rma request failed :request=%s,Error=%s"}
AssetPersistTestRequestRCVD	{"EFA-002085", "Asset persist test request received :request=%s"}
AssetPersistTestRequestSuccess{	"EFA-002086", "Asset persist test request success :request=%s"}
AssetPersistTestRequestFailed	{"EFA-002087", "Asset persist test request failed :request=%s,Error=%s"}
AssetChangeGlobalMTURequestRCVD	{"EFA-002088", "Asset change global mtu request received :request=%s"}
AssetChangeGlobalMTURequestSuccess	{"EFA-002089", "Asset change global mtu request success :request=%s"}
AssetChangeGlobalMTURequestFailed	{"EFA-002090", "Asset change global mtu request failed :request=%s,Error=%s"}
AssetUpdateDeviceDetailsRequestRCVD	{"EFA-002091", "Asset update device details request received :request=%s"}
AssetUpdateDeviceDetailsRequestSuccess	{"EFA-002092", "Asset update device details request success :request=%s"}
AssetUpdateDeviceDetailsRequestFailed	{"EFA-002093", "Asset update device details request failed :request=%s,Error=%s"}
AssetChangePortDampeningRequestFailed	{"EFA-002094", "Asset device state change request received :request=%s"}
AssetDeviceStateHandlerSuccess	{"EFA-002095", "Asset device state change request success :request=%s"}
AssetDeviceStateHandlerFailed	{"EFA-002096", "Asset device state change request failed :request=%s,Error=%s"}
AssetUpdateDeviceDetailsRequestFailed	{"EFA-002097", "Asset delete device state change request received :request=%s"}

XCO Notification Type	XCO Task
AssetRemovePortFecRequestRCVD	{"EFA-002098", "Asset delete device state change request success :request=%s"}
AssetDeleteDeviceStateRequestFailed	{"EFA-002099", "Asset delete device state change request failed :request=%s,Error=%s"}
AssetSnmpCommunityCreateRequestRCVD	{"EFA-002100", "Asset snmp community create request received :request=%s"}
AssetSnmpCommunityCreateRequestSuccess	{"EFA-002101", "Asset snmp community create request success :request=%s"}
AssetSnmpCommunityCreateRequestFailed	{"EFA-002102", "Asset snmp community create request failed :request=%s,Error=%s"}
AssetSnmpCommunityDeleteRequestRCVD	{"EFA-002103", "Asset snmp community delete request received :request=%s"}
AssetSnmpCommunityDeleteRequestSuccess	{"EFA-002104", "Asset snmp community delete request success :request=%s"}
AssetSnmpCommunityDeleteRequestFailed	{"EFA-002105", "Asset snmp community delete request failed :request=%s,Error=%s"}
AssetGetServiceApprovalForProcessRequestRCVD	{"EFA-002106", "Asset Get Service Approval For Process request received :request=%s"}
AssetGetServiceApprovalForProcessRequestRCVD	{"EFA-002107", "Asset Get Service Approval For Process request success :request=%s"}
AssetGetServiceApprovalForProcessRequestFailed	{"EFA-002108", "Asset Get Service Approval For Process request failed :request=%s,Error=%s"}
AssetChangePortFecRequestRCVD	{"EFA-002109", "Asset change port FEC request received :request=%s"}
AssetChangePortFecRequestSuccess	{"EFA-002110", "Asset change port FEC request success :request=%s"}
AssetChangePortFecRequestFailed	{"EFA-002111", "Asset change port FEC request failed :request=%s,Error=%s"}
AssetRemovePortFecRequestRCVD	{"EFA-002112", "Asset remove port FEC request received :request=%s"}
AssetRemovePortFecRequestSuccess	{"EFA-002113", "Asset remove port FEC request success :request=%s"}
AssetSnmpUserDeleteRequestFailed	{"EFA-002114", "Asset remove port FEC request failed :request=%s,Error=%s"}
AssetChangePortDampeningRequestRCVD	{"EFA-002115", "Asset change port dampening request received :request=%s"}
AssetChangePortFecRequestFailed	{"EFA-002116", "Asset change port dampening request success :request=%s"}
AssetChangePortDampeningRequestFailed	{"EFA-002117", "Asset change port dampening request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
AssetForceRemovalFwdlStatusRequestRCVD	{"EFA-002118", "Asset remove port dampening request received :request=%s"}
AssetRemovePortDampeningRequestSuccess	{"EFA-002119", "Asset remove port dampening request success :request=%s"}
AssetRemovePortDampeningRequestFailed	{"EFA-002120", "Asset remove port dampening request failed :request=%s,Error=%s"}
AssetNtpServerCreateRequestRCVD	{"EFA-002121", "Asset ntp server create request received :request=%s"}
AssetNtpServerCreateRequestSuccess	{"EFA-002122", "Asset ntp server create request success :request=%s"}
AssetNtpServerCreateRequestFailed	{"EFA-002123", "Asset ntp server create request failed :request=%s,Error=%s"}
AssetNtpServerDeleteRequestRCVD	{"EFA-002124", "Asset ntp server delete request received :request=%s"}
AssetNtpServerDeleteRequestSuccess	{"EFA-002125", "Asset ntp server delete request success :request=%s"}
AssetNtpServerCreateRequestRCVD	{"EFA-002126", "Asset ntp server delete request failed :request=%s,Error=%s"}
AssetGetServiceApprovalForProcessRequest RCVD	{"EFA-002127", "Asset disable ntp serve request received :request=%s"}
AssetNtpDisableRequestSuccess	{"EFA-002128", "Asset disable ntp serve request success :request=%s"}
AssetNtpDisableRequestSuccess	{"EFA-002129", "Asset disable ntp serve request failed :request=%s,Error=%s"}
AssetSnmpUserCreateRequestRCVD	{"EFA-002130", "Asset snmp user create request received :request=%s"}
AssetSnmpUserCreateRequestSuccess	{"EFA-002131", "Asset snmp user create request success :request=%s"}
AssetSnmpUserCreateRequestFailed	{"EFA-002132", "Asset snmp user create request failed :request=%s,Error=%s"}
AssetSnmpUserCreateRequestFailed	{"EFA-002133", "Asset snmp user delete request received :request=%s"}
AssetSnmpUserDeleteRequestSuccess	{"EFA-002134", "Asset snmp user delete request success :request=%s"}
AssetSnmpUserDeleteRequestFailed	{"EFA-002135", "Asset snmp user delete request failed :request=%s,Error=%s"}
AssetSnmpHostCreateRequestRCVD	{"EFA-002136", "Asset snmp host create request received :request=%s"}
AssetSnmpHostCreateRequestSuccess	{"EFA-002137", "Asset snmp host create request success :request=%s"}

XCO Notification Type	XCO Task
AssetSnmpHostDeleteRequestRCVD	{"EFA-002138", "Asset snmp host create request failed :request=%s,Error=%s"}
AssetSnmpHostDeleteRequestRCVD	{"EFA-002139", "Asset snmp host delete request received :request=%s"}
AssetSnmpHostDeleteRequestSuccess	{"EFA-002140", "Asset snmp host delete request success :request=%s"}
AssetSnmpHostDeleteRequestFailed	{"EFA-002141", "Asset snmp host delete request failed :request=%s,Error=%s"}
AssetChangeRedundantManagmentRequestRCVD	{"EFA-002142", "Asset change redundant managment request received :request=%s"}
AssetChangeRedundantManagmentRequestSuccess	{"EFA-002143", "Asset change redundant managment request success :request=%s"}
AssetChangeRedundantManagmentRequestFailed	{"EFA-002144", "Asset change redundant managment request failed :request=%s,Error=%s"}
AssetDebugDRCRequestRCVD	{"EFA-002145", "Asset debug drift-reconcile request received :request=%s"}
AssetDebugDRCRequestSuccess	{"EFA-002146", "Asset debug drift-reconcile request success :request=%s"}
AssetDebugDRCRequestFailed	{"EFA-002147", "Asset debug drift-reconcile request failed :request=%s,Error=%s"}
AssetForceRemovalFwdlStatusRequestRCVD	{"EFA-002148", "Asset force removal of fwdl status request received :request=%s"}
AssetForceRemovalFwdlStatusRequestSuccess	{"EFA-002149", "Asset force removal of fwdl status request success :request=%s"}
AssetForceRemovalFwdlStatusRequestFailed	{"EFA-002150", "Asset force removal of fwdl status request failed :request=%s,Error=%s"}
AssetGetFwdlInProgressDevicesRequestRCVD	{"EFA-002151", "Asset Get Fwdl In Progress Devices status request received :request=%s"}
AssetGetFwdlInProgressDevicesRequestSuccess	{"EFA-002152", "Asset Get Fwdl In Progress Devices status request success :request=%s"}
AssetGetFwdlInProgressDevicesRequestFailed	{"EFA-002153", "Asset Get Fwdl In Progress Devices status request failed :request=%s,Error=%s"}
AssetSetDeviceTimezoneRequestFailed	{"EFA-002154", "Asset set timezone request received :request=%s"}
AssetSetDeviceTimezoneRequestSuccess	{"EFA-002155", "Asset set timezone request success :request=%s"}
AssetSetDeviceTimezoneRequestFailed	{"EFA-002156", "Asset set timezone request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
AssetUnsetDeviceTimezoneRequestRCVD	{"EFA-002157", "Asset unset timezone request received :request=%s"}
AssetUnsetDeviceTimezoneRequestSuccess	{"EFA-002158", "Asset unset timezone request success :request=%s"}
AssetUnsetDeviceTimezoneRequestFailed	{"EFA-002159", "Asset unset timezone request failed :request=%s,Error=%s"}
AssetAddLocationRequestRCVD	{"EFA-002160", "Asset add location request received :request=%s"}
AssetAddLocationRequestSuccess	{"EFA-002161", "Asset add location request success :request=%s"}
AssetAddLocationRequestFailed	{"EFA-002162", "Asset add location request failed :request=%s,Error=%s"}
AssetGetAllLocationRequestRCVD	{"EFA-002163", "Asset get all location request received :request=%s,Error=%s"}
AssetGetAllLocationRequestSuccess	{"EFA-002164", "Asset get all location request success :request=%s,Error=%s"}
AssetGetAllLocationRequestFailed	{"EFA-002165", "Asset get all location request failed :request=%s,Error=%s"}
AssetDeleteLocationRequestRCVD	{"EFA-002166", "Asset delete location request received :request=%s"}
AssetDeleteLocationRequestSuccess	{"EFA-002167", "Asset delete location request success :request=%s"}
AssetDeleteLocationRequestFailed	{"EFA-002168", "Asset delete location request failed :request=%s,Error=%s"}
AssetUpdateLocationRequestRCVD	{"EFA-002169", "Asset update location request received :request=%s"}
AssetUpdateLocationRequestSuccess	{"EFA-002170", "Asset update location request success :request=%s"}
AssetUpdateLocationRequestFailed	{"EFA-002171", "Asset update location request failed :request=%s,Error=%s"}
AssetRestoreFinishedRequestRCVD	{"EFA-002172", "Asset restore finished request received :request=%s"}
AssetRestoreFinishedRequestSuccess	{"EFA-002173", "Asset restore finished request success :request=%s"}
AssetRestoreFinishedRequestFailed	{"EFA-002174", "Asset restore finished request failed :request=%s,Error=%s"}
AssetRemoveInterfaceMtuRequestRCVD	{"EFA-002175", "Asset unset interface mtu request received :request=%s"}
AssetRemoveInterfaceMtuRequestSuccess	{"EFA-002176", "Asset unset interface mtu request success :request=%s"}
AssetRemoveInterfaceMtuRequestFailed	{"EFA-002177", "Asset unset interface mtu request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
AssetReloadSwitchesRequestRCVD	{"EFA-002178", "Asset execute reload request received :request=%s"}
AssetReloadSwitchesRequestSuccess	{"EFA-002179", "Asset execute reload request success :request=%s"}
AssetReloadSwitchesRequestFailed	{"EFA-002180", "Asset execute reload request failed :request=%s,Error=%s"}
AssetClearAllRoutesRequestRCVD	{"EFA-002181", "Asset clear all routes request received :request=%s"}
AssetClearAllRoutesRequestSuccess	{"EFA-002182", "Asset clear all routes request success :request=%s"}
AssetClearAllRoutesRequestFailed	{"EFA-002183", "Asset clear all routes request failed :request=%s,Error=%s"}
AssetDriftAndReconcileSwitchesExecuteHandlerRCVD	{"EFA-002184", "Asset drift-reconcile switches execute request received :request=%s"}
AssetDriftAndReconcileSwitchesExecuteHandlerSuccess	{"EFA-002185", "Asset drift-reconcile switches execute request success :request=%s"}
AssetDriftAndReconcileSwitchesExecuteHandlerFailed	{"EFA-002186", "Asset drift-reconcile switches execute request failed :request=%s,Error=%s"}
AssetSnmpViewCreateRequestRCVD	{"EFA-002187", "Asset snmp view create request received :request=%s"}
AssetSnmpViewCreateRequestSuccess	{"EFA-002188", "Asset snmp view create request success :request=%s"}
AssetSnmpViewCreateRequestFailed	{"EFA-002189", "Asset snmp view create request failed :request=%s,Error=%s"}
AssetSnmpViewDeleteRequestRCVD	{"EFA-002190", "Asset snmp view delete request received :request=%s"}
AssetSnmpViewDeleteRequestSuccess	{"EFA-002191", "Asset snmp view delete request success :request=%s"}
AssetSnmpViewDeleteRequestRCVD	{"EFA-002192", "Asset snmp view delete request failed :request=%s,Error=%s"}
AssetSnmpViewGetRequestRCVD	{"EFA-002193", "Asset snmp view get request received :request=%s"}
AssetSnmpViewGetRequestSuccess	{"EFA-002194", "Asset snmp view get request success :request=%s"}
AssetSnmpViewGetRequestFailed	{"EFA-002195", "Asset snmp view get request failed :request=%s,Error=%s"}
AssetThresholdMonitorCreateRequestRCVD	{"EFA-002196", "Asset threshold monitor create request received :request=%s"}

XCO Notification Type	XCO Task
AssetThresholdMonitorCreateRequestSuccess	{"EFA-002197", "Asset threshold monitor create request success :request=%s"}
AssetThresholdMonitorCreateRequestFailed	{"EFA-002198", "Asset threshold monitor create request failed :request=%s,Error=%s"}
AssetThresholdMonitorDeleteRequestRCVD	{"EFA-002199", "Asset threshold monitor delete request received :request=%s"}
AssetThresholdMonitorDeleteRequestSuccess	{"EFA-002200", "Asset threshold monitor delete request success :request=%s"}
AssetThresholdMonitorDeleteRequestFailed	{"EFA-002201", "Asset threshold monitor delete request failed :request=%s,Error=%s"}
AssetThresholdMonitorGetRequestRCVD	{"EFA-002202", "Asset threshold monitor get request received :request=%s"}
AssetThresholdMonitorGetRequestSuccess{	{"EFA-002203", "Asset threshold monitor get request success :request=%s"}
AssetThresholdMonitorGetRequestFailed	{"EFA-002204", "Asset threshold monitor get request failed :request=%s,Error=%s"}
AssetSecureSettingsEnableRequestRCVD	{"EFA-002205", "Asset secure settings enable request received :request=%s"}
AssetSecureSettingsEnableRequestSuccess	{"EFA-002206", "Asset secure settings enable request success :request=%s"}
AssetSecureSettingsEnableRequestFailed	{"EFA-002207", "Asset secure settings enable request failed :request=%s,Error=%s"}
AssetSecureSettingsDisableRequestRCVD	{"EFA-002208", "Asset secure settings disable request received :request=%s"}
AssetSecureSettingsDisableRequestSuccess	{"EFA-002209", "Asset secure settings disable request success :request=%s"}
AssetSecureSettingsDisableRequestFailed	{"EFA-002210", "Asset secure settings disable request failed :request=%s,Error=%s"}
AssetSecureSettingsResetRequestRCVD	{"EFA-002211", "Asset secure settings reset request received :request=%s"}
AssetSecureSettingsResetRequestSuccess	{"EFA-002212", "Asset secure settings reset request success :request=%s"}
AssetSecureSettingsResetRequestFailed	{"EFA-002213", "Asset secure settings reset request failed :request=%s,Error=%s"}
AssetSecureSettingsUpdateRequestRCVD	{"EFA-002214", "Asset secure settings update request received :request=%s"}
AssetSecureSettingsUpdateRequestSuccess	{"EFA-002215", "Asset secure settings update request success :request=%s"}
AssetSecureSettingsUpdateRequestFailed	{"EFA-002216", "Asset secure settings update request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
AssetSecureSettingsApplyRequestRCVD	{"EFA-002217", "Asset secure settings apply request received :request=%s"}
AssetSecureSettingsApplyRequestSuccess{	{"EFA-002218", "Asset secure settings apply request success :request=%s"}
AssetSecureSettingsApplyRequestFailed	{"EFA-002219", "Asset secure settings apply request failed :request=%s,Error=%s"}
AssetSecureSettingsGetRequestRCVD	{"EFA-002220", "Asset secure settings get request received :request=%s"}
AssetSecureSettingsGetRequestSuccess	{"EFA-002221", "Asset secure settings get request success :request=%s"}
AssetSecureSettingsGetRequestFailed	{"EFA-002222", "Asset secure settings get request failed :request=%s,Error=%s"}
AssetSnmpGroupUpdateRequestRCVD	{"EFA-002223", "Asset snmp group update request received :request=%s"}
AssetSnmpGroupUpdateRequestSuccess	{"EFA-002224", "Asset snmp group update request success :request=%s"}
AssetSnmpGroupUpdateRequestFailed	{"EFA-002225", "Asset snmp group update request failed :request=%s,Error=%s"}
AssetSnmpGroupGetRequestRCVD	{"EFA-002226", "Asset snmp group get request received :request=%s"}
AssetSnmpGroupGetRequestSuccess	{"EFA-002227", "Asset snmp group get request success :request=%s"}
AssetSnmpGroupGetRequestFailed	{"EFA-002228", "Asset snmp group get request failed :request=%s,Error=%s"}
AssetSnmpUseVrfCreateRequestRCVD	{"EFA-002229", "Asset snmp use-vrf create request received :request=%s"}
AssetSnmpUseVrfCreateRequestSuccess	{"EFA-002230", "Asset snmp use-vrf create request success :request=%s"}
AssetSnmpUseVrfCreateRequestFailed	{"EFA-002231", "Asset snmp use-vrf create request failed :request=%s,Error=%s"}
AssetSnmpUseVrfDeleteRequestRCVD	{"EFA-002232", "Asset snmp use-vrf delete request received :request=%s"}
AssetSnmpUseVrfDeleteRequestSuccess	{"EFA-002233", "Asset snmp use-vrf delete request success :request=%s"}
AssetSnmpUseVrfDeleteRequestFailed	{"EFA-002234", "Asset snmp use-vrf delete request failed :request=%s,Error=%s"}
AssetSnmpUseVrfGetRequestRCVD	{"EFA-002235", "Asset snmp use-vrf get request received :request=%s"}

XCO Notification Type	XCO Task
AssetSnmpUseVrfGetRequestSuccess	{"EFA-002236", "Asset snmp use-vrf get request success :request=%s"}
AssetSnmpUseVrfGetRequestFailed	{"EFA-002237", "Asset snmp use-vrf get request failed :request=%s,Error=%s"}

Policy Service Notifications

Note
%S is policy name.

XCO Notification Type	XCO Task
PolicyPrefixListCreateRequestRCVD	{"EFA-003001", "Policy prefix-list create request received :request=%s"}
PolicyPrefixListCreateRequestSuccess	{"EFA-003002", "Policy prefix-list create request success :request=%s"}
PolicyPrefixListCreateRequestFailed	{"EFA-003003", "Policy prefix-list create request failed :request=%s,Error=%s"}
PolicyPrefixListUpdateRequestRCVD	{"EFA-003004", "Policy prefix-list update request received :request=%s"}
PolicyPrefixListUpdateRequestSuccess	{"EFA-003005", "Policy prefix-list update request success :request=%s"}
PolicyPrefixListUpdateRequestFailed	{"EFA-003006", "Policy prefix-list update request failed :request=%s,Error=%s"}
PolicyPrefixListDeleteRequestRCVD	{"EFA-003007", "Policy prefix-list delete request received :request=%s"}
PolicyPrefixListDeleteRequestSuccess	{"EFA-003008", "Policy prefix-list delete request success :request=%s"}
PolicyPrefixListDeleteRequestFailed	{"EFA-003009", "Policy prefix-list delete request failed :request=%s,Error=%s"}
PolicyRouteMapCreateRequestRCVD	{"EFA-003010", "Policy RouteMap create request received :request=%s"}
PolicyRouteMapCreateRequestSuccess	{"EFA-003011", "Policy RouteMap create request success :request=%s"}
PolicyRouteMapCreateRequestFailed	{"EFA-003012", "Policy RouteMap create request failed :request=%s,Error=%s"}
PolicyRouteMapUpdateRequestRCVD	{"EFA-003013", "Policy RouteMap update request received :request=%s"}
PolicyRouteMapUpdateRequestSuccess	{"EFA-003014", "Policy RouteMap update request success :request=%s"}
PolicyRouteMapUpdateRequestFailed	{"EFA-003015", "Policy RouteMap update request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
PolicyRouteMapDeleteRequestRCVD	{"EFA-003016", "Policy RouteMap delete request received :request=%s"}
PolicyRouteMapDeleteRequestSuccess	{"EFA-003017", "Policy RouteMap delete request success :request=%s"}
PolicyRouteMapDeleteRequestFailed	{"EFA-003018", "Policy RouteMap delete request failed :request=%s,Error=%s"}
PolicyRouteMapMatchCreateRequestRCVD	{"EFA-003019", "Policy RouteMap match create request received :request=%s"}
PolicyRouteMapMatchCreateRequestSuccess	{"EFA-003020", "Policy RouteMap match create request success :request=%s"}
PolicyRouteMapMatchCreateRequestFailed	{"EFA-003021", "Policy RouteMap match create request failed :request=%s,Error=%s"}
PolicyRouteMapMatchDeleteRequestRCVD	{"EFA-003022", "Policy RouteMap match delete request received :request=%s"}
PolicyRouteMapMatchDeleteRequestSuccess	{"EFA-003023", "Policy RouteMap match delete request success :request=%s"}
PolicyRouteMapMatchDeleteRequestFailed	{"EFA-003024", "Policy RouteMap match delete request failed :request=%s,Error=%s"}
PolicyCommunityListCreateRequestRCVD	{"EFA-003025", "Policy community-list create request received :request=%s"}
PolicyCommunityListCreateRequestSuccess	{"EFA-003026", "Policy community-list create request success :request=%s"}
PolicyCommunityListCreateRequestFailed	{"EFA-003027", "Policy community-list create request failed :request=%s,Error=%s"}
PolicyCommunityListUpdateRequestRCVD	{"EFA-003028", "Policy community-list update request received :request=%s"}
PolicyCommunityListUpdateRequestSuccess	{"EFA-003029", "Policy community-list update request success :request=%s"}
PolicyCommunityListUpdateRequestFailed	{"EFA-003030", "Policy community-list update request failed :request=%s,Error=%s"}
PolicyCommunityListDeleteRequestRCVD	{"EFA-003031", "Policy community-list delete request received :request=%s"}
PolicyCommunityListDeleteRequestSuccess	{"EFA-003032", "Policy community-list delete request success :request=%s"}
PolicyCommunityListDeleteRequestFailed	{"EFA-003033", "Policy community-list delete request failed :request=%s,Error=%s"}
PolicyExtCommunityListCreateRequestRCVD	{"EFA-003034", "Policy ExtCommunityList create request received :request=%s"}
PolicyExtCommunityListCreateRequestSuccess	{"EFA-003035", "Policy ExtCommunityList create request success :request=%s"}

XCO Notification Type	XCO Task
PolicyExtCommunityListCreateRequestFailed	{"EFA-003036", "Policy ExtCommunityList create request failed :request=%s,Error=%s"}
PolicyExtCommunityListUpdateRequestRCVD	{"EFA-003037", "Policy ExtCommunityList update request received :request=%s"}
PolicyExtCommunityListUpdateRequestSuccess	{"EFA-003038", "Policy ExtCommunityList update request success :request=%s"}
PolicyExtCommunityListUpdateRequestFailed	{"EFA-003039", "Policy ExtCommunityList update request failed :request=%s,Error=%s"}
PolicyExtCommunityListDeleteRequestRCVD	{"EFA-003040", "Policy ExtCommunityList delete request received :request=%s"}
PolicyExtCommunityListDeleteRequestSuccess	{"EFA-003041", "Policy ExtCommunityList delete request success :request=%s"}
PolicyExtCommunityListDeleteRequestFailed	{"EFA-003042", "Policy ExtCommunityList delete request failed :request=%s,Error=%s"}
PolicyRouteMapSetCreateRequestRCVD	{"EFA-003043", "Policy RouteMap set create request received :request=%s"}
PolicyRouteMapSetCreateRequestSuccess	{"EFA-003044", "Policy RouteMap set create request success :request=%s"}
PolicyRouteMapSetCreateRequestFailed	{"EFA-003045", "Policy RouteMap set create request failed :request=%s,Error=%s"}
PolicyRouteMapSetDeleteRequestRCVD	{"EFA-003046", "Policy RouteMap set delete request received :request=%s"}
PolicyRouteMapSetDeleteRequestSuccess	{"EFA-003047", "Policy RouteMap set delete request success :request=%s"}
PolicyRouteMapSetDeleteRequestFailed	{"EFA-003048", "Policy RouteMap set delete request failed :request=%s,Error=%s"}
PolicyQosMapCreateRequestRCVD	{"EFA-003049", "Policy qos-map create request received :request=%s"}
PolicyQosMapCreateRequestSuccess	{"EFA-003050", "Policy qos-map create request success :request=%s"}
PolicyQosMapCreateRequestFailed	{"EFA-003051", "Policy qos-map create request failed :request=%s,Error=%s"}
PolicyQosMapUpdateRequestRCVD	{"EFA-003052", "Policy qos-map update request received :request=%s"}
PolicyQosMapUpdateRequestSuccess	{"EFA-003053", "Policy qos-map update request success :request=%s"}
PolicyQosMapUpdateRequestFailed	{"EFA-003054", "Policy qos-map update request failed :request=%s,Error=%s"}
PolicyQosMapDeleteRequestRCVD	{"EFA-003055", "Policy qos-map delete request received :request=%s"}

XCO Notification Type	XCO Task
PolicyQosMapDeleteRequestSuccess	{"EFA-003056", "Policy qos-map delete request success :request=%s"}
PolicyQosMapDeleteRequestFailed	{"EFA-003057", "Policy qos-map delete request failed :request=%s,Error=%s"}
PolicyQosProfileCreateRequestRCVD	{"EFA-003058", "Policy qos-profile create request received :request=%s"}
PolicyQosProfileCreateRequestSuccess	{"EFA-003059", "Policy qos-profile create request success :request=%s"}
PolicyQosProfileCreateRequestFailed	{"EFA-003060", "Policy qos-profile create request failed :request=%s,Error=%s"}
PolicyQosProfileUpdateRequestRCVD	{"EFA-003061", "Policy qos-profile update request received :request=%s"}
PolicyQosProfileUpdateRequestSuccess	{"EFA-003062", "Policy qos-profile update request success :request=%s"}
PolicyQosProfileUpdateRequestFailed	{"EFA-003063", "Policy qos-profile update request failed :request=%s,Error=%s"}
PolicyQosProfileDeleteRequestRCVD	{"EFA-003064", "Policy qos-profile delete request received :request=%s"}
PolicyQosProfileDeleteRequestSuccess	{"EFA-003065", "Policy qos-profile delete request success :request=%s"}
PolicyQosProfileDeleteRequestFailed	{"EFA-003066", "Policy qos-profile delete request failed :request=%s,Error=%s"}
PolicyQosServiceMapCreateRequestRCVD	{"EFA-003067", "Policy qos-service-map create request received :request=%s"}
PolicyQosServiceMapCreateRequestSuccess	{"EFA-003068", "Policy qos-service-map create request success :request=%s"}
PolicyQosServiceMapCreateRequestFailed	{"EFA-003069", "Policy qos-service-map create request failed :request=%s,Error=%s"}
PolicyQosServiceMapUpdateRequestRCVD	{"EFA-003070", "Policy qos-service-map update request received :request=%s"}
PolicyQosServiceMapUpdateRequestSuccess	{"EFA-003071", "Policy qos-service-map update request success :request=%s"}
PolicyQosServiceMapUpdateRequestFailed	{"EFA-003072", "Policy qos-service-map update request failed :request=%s,Error=%s"}
PolicyQosServiceMapDeleteRequestRCVD	{"EFA-003073", "Policy qos-service-map delete request received :request=%s"}
PolicyQosServiceMapDeleteRequestSuccess	{"EFA-003074", "Policy qos-service-map delete request success :request=%s"}
PolicyQosServiceMapDeleteRequestFailed	{"EFA-003075", "Policy qos-service-map delete request failed :request=%s,Error=%s"}

XCO Notification Type	XCO Task
AssetSetInterfaceDescriptionRequestRCVD	{"EFA-003076", "Asset set interface description request received :request=%s"}
AssetSetInterfaceDescriptionRequestSuccess	{"EFA-003077", "Asset set interface description request success :request=%s"}
AssetSetInterfaceDescriptionRequestFailed	{"EFA-003078", "Asset set interface description request failed :request=%s,Error=%s"}
AssetRemoveInterfaceDescriptionRequestRCVD	{"EFA-003079", "Asset unset interface description request received :request=%s"}
AssetRemoveInterfaceDescriptionRequestSuccess	{"EFA-003080", "Asset unset interface description success :request=%s"}
AssetRemoveInterfaceDescriptionRequestFailed	{"EFA-003081", "Asset unset interface description request failed :request=%s,Error=%s"}

Notifications Service

Note
%S is a notification name.

XCO Notification Type	XCO Task
NotificationServiceHandlerConnectivityTestMsg	{"EFA-008000", "Test payload to verify %s connection"}

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 and later.

The following table summarizes the supported key-value pairs:

Payload Fields	APP_EVENTS (Task Events)	DEVICE_EVENTS (SLX Raslog Events)	HOST_EVENTS	APP_ALERTS (Fault-Alerts Only)	APP_ALARMS (Fault Alarms Only)
type	✓ (Task)	✓ (Event)	✓ (Event)	✓ (Alert)	✓ (Alarm)
timestamp	✓	✓	✓	✓	✓
severity	✓	✓	✓	✓	✓
message	✓	✓	✓	✓	✓
application	✓	✓	✓	✓	✓
source_ip	□	✓	✓	✓	✓

Payload Fields	APP_EVENTS (Task Events)	DEVICE_EVENTS (SLX Raslog Events)	HOST_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)*

Field	SD-ID (Structured Data ID)	Example	Description
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*
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*
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.
```

HOST Events RFC-5424 Format

This provides the common fields of the HOST_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*	*10.32.85.29*	IP Of XCO Node
App Name	*N/A*	*auditd*	Application generating syslog alerts. Possible values
Proc ID	*N/A*	*AUDITD_EVENTS*	*Process ID*
Msg ID	*N/A*	*Auditd Message ID*	*Raslog ID of the device event*
Structured Data	*meta*		
Message	*N/A*		*Actual auditd message*

Map HOST_EVENTS to RELP/Syslog fields (RFC-5424)

```
<14>1 2025-03-18T10:30:58.110734Z 10.32.85.28 audit 48734 AUDITD_EVENTS -
node=127.0.1.1 type=USER_START msg=audit(1742293858.104:1365): pid=48734 uid=0 auid=1000
ses=16 subj=unconfined msg='op=PAM:session_open
grantors=pam_selinux,pam_loginuid,pam_keyinit,pam_permit,pam_umask,pam_unix,pam_systemd,pa
m_mail,pam_limits,pam_env,pam_env,pam_selinux acct="user" exe="/usr/sbin/sshd"
hostname=134.141.202.209 addr=134.141.202.209 terminal=ssh res=success' UID="root"
AUID="user".
```

RELP (Syslog) and Webhook Remote Server Secure Notifications Configuration

Use this topic to learn about configuring Remove Syslog (RELP) and Webhook securely.

- Configure Remote Syslog (RELP) Securely on page 861
- Configure Webhook Securely on page 862

Configure Remote Syslog (RELP) Securely

You can configure remote Syslog (RELP) securely.

About This Task

Follow this procedure to configure remote Syslog (RELP) securely.

Procedure

1. Install certtool on the Ubuntu system.

```
sudo apt-get install gnutls-bin
```

2. Extract `relp-certs.tar.gz` and edit `ca.cfg` and `server.cfg` per guidelines (`client.cfg` is not needed if you are configuring the client from XCO).

Note

Download the tar file `relp-certs.tar .gz`.

For guidelines on updating `ca.cfg` and the `server.cfg` files, see Configuration Files Modification Guidelines on page 863.

3. Run the **`./1-generate-ca.sh`** command to generate CA files (ca.pem and ca-key.pem).
4. Run the **`generate-server.sh`** command to generate server certificates (server-cert.pem and server-key.pem).
5. Copy `extreme-ca-chain.pem` from EFA's `/apps/efadata/certs/ca` location to syslog server's `/opt/notification/` location (create folder if absent).

Note

If a third-party or CMPv2 client certificate is installed on XCO, copy the third-party or CMPv2 CA chain that signed the certificate instead of the `extreme-ca-chain.pem` file.

6. Copy `server-cert.pem` and `server-key.pem` to `/opt/notification/`.
7. Edit `/etc/rsyslog.conf` with root/sudo permission. Add the following lines above GLOBAL DIRECTIVES in the file:

```
module(load="imrelp")

input(type="imrelp" port="20514" tls="on" tls.caCert="/opt/notification/extreme-ca-chain.pem" tls.myCert="/opt/notification/server-cert.pem" tls.myPrivKey="/opt/notification/server-key.pem" tls.authMode="name" tls.permittedpeer=["tpvm128","10.20.63.128","10.20.48.245","10.20.246.7","10.20.241.7","10.20.63.70","tpvm-72","administrator-00","efa.extremenetworks.com","10.20.54.90"])
```

Note

- Ensure to include the XCO VIP address or the IP address/server address (where XCO is installed) in the list of IP addresses.
- If you used a third-party or CMPv2 CA chain in Step 5, set `tls.caCert` to that file instead of `extreme-cachain.pem`.

8. Restart rsyslog service.

```
sudo systemctl restart rsyslog
```

9. Complete the following RELP configuration in XCO:
   a. Copy `ca.pem` from syslog server to XCO's `/home/extreme`.
   b. Add subscriber.

```
efa notification subscribers add-syslog-relp --address=<syslog-IP>:20514 --
cacert=ca.pem --conn-timeout=30 (for example, 10.21.88.25:20514)
```

Note
Address is the Notification Server IP address.

   c. Run some of the XCO commands to verify syslog receipt (if the messages are forwarded).

   On XCO:

```
(efa:xmcdev)xmcdev@xmcdev-virtual-machine1:/opt/checkouts/efa/efa$ efa tenant
create --name=tnt3

Tenant created successfully.

--- Time Elapsed: 209.27997ms ---

<5>Apr  6 01:49:35 10.133.131.162(10.133.131.162)
{"application":"ts","device_ip":"","message":"Tenant create request
received :request={\"name\":\"tnt3\"}","scope":"user","status":"started","task":"EFA
-001001","timestamp":"2022-04-06T01:49:04+05:30","type":"Task","username":"xmcdev"}

<5>Apr  6 01:49:35 10.133.131.162(10.133.131.162)
{"application":"ts","device_ip":"","message":"Tenant create request
success :request={\"name\":\"tnt3\"}","scope":"user","status":"succeeded","task":"EF
A-001002","timestamp":"2022-04-06T01:49:04+05:30","type":"Task","username":"xmcdev"}
```

Configure Webhook Securely

You can configure Webhook securely.

About This Task

Follow this procedure to configure Webhook securely.

Procedure

1. Install certtool on the Ubuntu system.

```
sudo apt-get install gnutls-bin
```

2. Extract `relp-certs.tar.gz` and edit `ca.cfg` and `server.cfg` per guidelines (`client.cfg` is not needed if you are configuring the client from XCO).

Note
Download the tar file `relp-certs.tar .gz`.

For guidelines on updating `ca.cfg` and the `server.cfg` files, see Configuration Files Modification Guidelines on page 863.

3. Run the **`./1-generate-ca.sh`** command to generate CA files (`ca.pem` and `cakey. pem`).
4. Run the **`generate-server.sh`** command to generate server certificates (`servercert.pem` and `server-key.pem`).

5. Copy `server-cert.pem` and `server-key.pem` to webhook notification server's `/home/user/subscriber` folder.
6. Edit `run.sh` file. Set `CERTS_DIR=~/subscriber` (the folder subscriber).

   Note
   If the server cert file is named `server-crt.pem` instead of `server-cert.pem` in the `run.sh` file, rename it to `server-crt.pem` in the subscriber folder to match.

7. Start: **`bash run.sh insecure=false`**.
8. Complete the following configuration in XCO:
   a. Copy `ca.pem` from webhook server to XCO.
   b. Add subscriber.

```
efa notification subscribers add-https --cacert=ca.pem --insecure=false
--urlhttps://10.20.61.58:5000--username=jarvis
--password=vision
```

   c. Run some of the XCO commands to verify if the messages are forwarded in webook server.

Configuration Files Modification Guidelines

Learn modification guidelines for `ca.cfg` and `server.cfg` files.

ca.cfg

Uncomment or update the following values:

- `organization = "Extreme Networks"`
- `unit = "Extreme"`
- `country = US`
- `cn = "Extreme CA"`
- `expiration_days = 3650`
- `challenge_password =` (leave empty)
- `ca`
- `signing_key`
- `cert_signing_key`

server.cfg

Uncomment or update the following values:

- `organization = "Extreme Networks"`
- `unit = "Extreme"`
- `cn = "ana-sowj.extremenetworks.com"`
- `expiration_days = 3650`

- `ip_address = "10.21.88.25"`
- `challenge_password` = (leave empty)
- `tls_www_client`
- `tls_www_server`

Configure Notification Subscriber With FQDN

You can configure a notification subscriber using a Fully Qualified Domain Name (FQDN).

About This Task

Follow this procedure to configure FQDN-based notification subscribers on TPVM with XCO.

Note

Server deployment requires a similar DNS configuration.

Procedure

1. Configure DNS servers in TPVM.

   Enter TPVM configuration mode and specify primary and secondary DNS servers for FQDN resolution.

   ```
   SW_93_TPVM_91(config-tpvm-TPVM)#dns primary-server 134.141.79.201 secondary-server
   134.141.79.202
   ```

   Verify configuration:

   ```
   SW_93_TPVM_91# show run tpvm
   tpvm TPVM
    auto-boot
    ntp 10.20.55.200
    ntp 10.31.2.80
    ntp 10.38.99.47
   ```

2. Add Syslog RELP subscriber with FQDN.

   Run the following command to add a notification subscriber targeting the FQDN endpoint:

   ```
   (efa:extreme)extreme@tpvm2:/etc/network$ efa notification subscribers add-syslog-relp --address
   syslog171.extremenetworks.com:20514 -insecure

           Successfully registered subscriber
   +-----------+-------------------------------------------------------------------------------------+
   | attribute | value                                                                               |
   +-----------+-------------------------------------------------------------------------------------+
   | id        | 5                                                                                   |
   | handler   | relp                                                                                |
   | endpoint  | syslog171.extremenetworks.com:20514                                                 |
   | config    | {"cacert":"","conn-timeout":10,"device-event":[],"filters":[],"host-auditd-key":[],
   |
   |           | "host-auditd-record-type":[],"host-event":[],"insecure":true,"minimum-severity":"",
   |
   |           | "rfc5424":false}                                                                    |
   +-----------+-------------------------------------------------------------------------------------+
   ```

3. Verify syslog reception.

   Messages successfully received at syslog server (syslog171.extremenetworks.com:20514):

```
$ efa fabric show

<13>1 2025-08-21T12:36:05.290735+00:00 10.20.54.92 - - - -
{"type":"Task","timestamp":"2025-08-21T05:36:05-07:00","severity":"info","message":"Fab
ric show executions request received :request={\"detail\":
[\"false\"]}","application":"fabric","source_ip":"","device_ip":"","username":"extreme"
,"message_id":"","hostname":"","logtype":"","task":"EFA-000037","scope":"user","status"
:"started","sequence_id":0,"alert_id":0,"alarm_id":0,"resource":"","alarm_type":"","ala
rm_cause":"","alert_data":null,"alarm_data":null}

<13>1 2025-08-21T12:36:05.298383+00:00 10.20.54.92 - - - -
{"type":"Task","timestamp":"2025-08-21T05:36:05-07:00","severity":"info","message":"Fab
ric show executions request success :request={\"detail\":
[\"false\"]}","application":"fabric","source_ip":"","device_ip":"","username":"extreme"
,"message_id":"","hostname":"","logtype":"","task":"EFA-000038","scope":"user","status"
:"succeeded","sequence_id":0,"alert_id":0,"alarm_id":0,"resource":"","alarm_type":"","a
larm_cause":"","alert_data":null,"alarm_data":null}
```

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, 4.1.0*.

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 143.

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 23: 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 (SLX), 1-65535 (Extreme OS ONE),

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.
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	

Field	Identify Drift	Reconcile config	Idempotency	Comments
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.
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	

Field	Identify Drift	Reconcile config	Idempotency	Comments
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	

Configure Device SNMP Use-VRF

An SNMP Use-vrf is a group of VRFs for SNMP listening service.

About This Task

Follow this procedure to configure device SNMP Use-VRF.

For information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Table 24: Drift Reconcile & Idempotency support for Device SNMP Use-VRF

Identify Drift	Reconcile configuration	Idempotency
Yes	Yes	No

Procedure

1. To configure an SNMP use-vrf, run the **`efa inventory device snmp use-vrf create`** command.

   The following example creates a use-vrf on specified devices:

```
efa inventory device snmp use-vrf create --name vrf1 --ip 10.139.44.159-160 --shutdown
+---------------+---------+----------+---------+--------+
|  IP Address   | Use VRF | Shutdown | Status  | Reason |
+---------------+---------+----------+---------+--------+
| 10.139.44.159 | vrf1    | true     | Success |        |
+---------------+---------+----------+---------+--------+
| 10.139.44.160 | vrf1    | true     | Success |        |
+---------------+---------+----------+---------+--------+
Snmp use-vrf details
```

2. To delete an SNMP Use-vrf, run the **`efa inventory device snmp use-vrf delete`** command.

   The following example deletes a use-vrf on specified devices:

```
efa inventory device snmp use-vrf delete --name vrf1 --ip 10.139.44.159-160
+---------------+-----------+---------+--------+
```

```
|  IP Address   | User Name | Status  | Reason |
+---------------+-----------+---------+--------+
| 10.139.44.160 | vrf1      | Success |        |
+---------------+-----------+---------+--------+
| 10.139.44.159 | vrf1      | Success |        |
+---------------+-----------+---------+--------+
Snmp user details
```

3. To list an SNMP Use-vrfs on device, run the **`efa inventory device snmp use-vrf list`** command.

   The following example shows the SNMP use-vrfs for the specified devices:

```
efa inventory device snmp use-vrf list --ip 10.139.44.159-160
+---------------+-------------+----------+-----------------+
|  IP Address   |   Use VRF   | Shutdown |    AppState     |
+---------------+-------------+----------+-----------------+
| 10.139.44.159 | default-vrf | false    | cfg-not-managed |
+---------------+-------------+----------+-----------------+
| 10.139.44.159 | mgmt-vrf    | false    | cfg-not-managed |
+---------------+-------------+----------+-----------------+
| 10.139.44.159 | vrf1        | true     | cfg-in-sync     |
+---------------+-------------+----------+-----------------+
| 10.139.44.160 | default-vrf | false    | cfg-not-managed |
+---------------+-------------+----------+-----------------+
| 10.139.44.160 | mgmt-vrf    | false    | cfg-not-managed |
+---------------+-------------+----------+-----------------+
| 10.139.44.160 | vrf1        | true     | cfg-in-sync     |
+---------------+-------------+----------+-----------------+
Snmp use-vrf details
```

Configure Device SNMP Group

You can update or delete a device SNMP group.

About This Task

Follow this procedure to update or delete a device SNMP group.

- The SNMP v2c/v3 group is created automatically when you create your first associated community or user.
- The SNMP v2c/v3 group is deleted automatically when you delete the last associated community or user.

For information about commands and supported parameters, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

1. Update an SNMP group.
   a. Run the **`efa inventory device snmp group update`** command on SLX devices.

```
efa inventory device snmp group update --ip 10.64.160.30 --name group1 --version
v2c --read-view view1 --notify-view view1
+--------------+------------+---------+-----------+------------+-------------
+------------+---------+---------+
|  IP Address  | Group Name | Version | Read view
| Write view | Notify view | Auth Level | Status  | Reason |
+--------------+------------+---------+-----------+------------+-------------
+------------+---------+---------+
| 10.64.160.30 | group1     | v2c     | view1     |
| view1        |            | Success |        |
```

```
+--------------+------------+---------+-----------+------------+-------------
+------------+---------+--------+
Snmp group details
--- Time Elapsed: 10.824500691s ---
```

b. Run the **show running-config snmp-server** command on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf mgmt-
vrf


snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c read view1 notify
view1

snmp-server group group1 v3 priv
```

2. To list SNMP groups on the device, run the **efa inventory device snmp group list** command.

   The following example lists groups on a specific device:

```
efa inventory device snmp group list --ip 10.139.44.159
+---------------+-------------+---------+-------------+------------+-------------
+------------+-----------------+
|  IP Address   | Group Name  | Version |  Read view  | Write
view | Notify view | Auth Level |    AppState     |
+---------------+-------------+---------+-------------+------------+-------------
+------------+-----------------+
| 10.139.44.159 | efav3Group  | v3      |             |
| efav3View   | noauth     | cfg-not-managed |
+---------------+-------------+---------+-------------+------------+-------------
+------------+-----------------+
| 10.139.44.159 | test-grp-v2 | v2c     | test-view-1 | efav3View
|             | noauth     | cfg-in-sync     |
+---------------+-------------+---------+-------------+------------+-------------
+------------+-----------------+
Snmp group details
```

SNMP Trap & Inform Delivery

Learn about forwarding SNMP traps and informs from your managed switches to external SNMP managers.

XCO gives you two independent ways to deliver SNMP traps and informs from your managed switches to an external SNMP manager, also called a Network Management System (NMS). You can use either path on its own, or run both at the same time.

XCO Proxy Path

Your switches send their notifications to XCO, and XCO forwards them on to every NMS you have registered. This path is configured for you automatically the moment a switch is registered in XCO — there is nothing to set up on the switch. To start receiving notifications, you simply register your NMS as a subscriber with a single API call.

Direct Path

You configure a switch to send its notifications straight to your NMS, without XCO sitting in the middle. This path takes a few more steps per switch, but it supports reliable SNMPv3 informs with acknowledgment and retry, and it lets you control exactly which switch sends what, and to where.

Note
Throughout this topic, replace <XCO_IP> with the IP address of your XCO instance and <token> with a valid API authorization token. Example IP addresses such as 10.0.0.50 and 10.1.1.1 are placeholders — substitute your own.

Supported Features

- **Two delivery paths:** a zero-touch XCO Proxy path and a fully controllable Direct path — usable separately or together.
- **Zero-touch proxy setup:** The proxy path is auto-configured when a device is registered, so administrators do nothing on the switch.
- **Reliable informs:** The Direct path supports SNMPv3 informs with acknowledgment and retry for dependable delivery.
- **Broad platform support:** Works with both SLX-OS and OS ONE managed switches.
- **Up to four proxy subscribers:** Register as many as four external SNMP managers on the proxy path.

Choosing How to Deliver Notifications

XCO supports three deployment models. Pick the one that matches how much control and reliability you need.

Model	How notifications flow	Setup effort	Reliability	Informs?
A — XCO Proxy (Default)	Switch → XCO → NMS	Zero-touch, plus one API call	Best-effort	No
B — Direct to External Manager	Switch → NMS	2–5 API calls per switch	Configurable	Yes (v3)
C — Hybrid (Both)	Both paths at once	Both sets of calls	Both	Partial

Which Model Should I Use?

If you need...	Use
The fastest possible setup, and occasional trap loss is acceptable	**Model A (Proxy)**
Reliable delivery with inform acknowledgments	**Model B (Direct)**
Unmediated monitoring for a compliance requirement	**Model B (Direct)**
To remove XCO as a single point of failure for SNMP	**Model B (Direct)**
Per-switch control over which notifications go where	**Model B (Direct)**
More than four external SNMP managers	**Model B (Direct)**
Central aggregation for dashboards and a reliable feed to a compliance NMS	**Model C (Hybrid)**

Best Practices

Prefer SNMPv3 over v2c wherever your NMS supports it. SNMPv3 authenticates and encrypts your notifications, while v2c sends a community string in the clear.

- **Use informs for anything you cannot afford to miss.** Informs are acknowledged and retried; traps are not. On SLX-OS, informs require SNMPv3.
- **Reuse XCO's default SNMP objects for a quick start,** or create your own for stricter isolation between the proxy and direct paths.
- **Stay within platform limits.** SLX-OS allows 6 hosts per version; OS ONE allows 18 total. The proxy path allows 4 subscribers.
- **Order a slow NMS last on the proxy path.** Proxy forwarding is sequential, so an unresponsive subscriber delays delivery to those after it.
- **Configure a VRF:** Only when your NMS is reachable exclusively through a non-default routing instance.

Before You Begin

To use either path (XCO proxy or Direct) you have the following:

- An XCO 4.1.0 instance with your switches registered.
- An external SNMP manager (NMS) reachable from either XCO (proxy path) or the switch (direct path) on UDP port 162.
- The SNMP credentials your NMS expects — a community string for SNMPv2c, or a username with authentication and privacy passphrases for SNMPv3.
- For SNMPv3 informs: your NMS's SNMP engine ID.

Platform & Version Support

Supported SNMP Versions

Version	SLX-OS	OS ONE
v1	Not supported	Not supported
v2c	Supported	Supported
v3	Supported	Supported

Configuration Objects by Platform

Object	Purpose	SLX-OS	OS ONE
SNMP View	OID tree include/exclude filter	Required	Required
SNMP Group	Access control; enables notify access	Required	Required
SNMP Community	v2c authentication string	Group required	Group required
SNMP User	v3 user with auth/priv credentials	Group required	Group required
SNMP Host	Trap/inform destination	Max 6 per version	Max 18 total
SNMP Use-VRF	VRF for SNMP traffic routing	Yes	Yes

Host Features by Platform

Feature	SLX-OS	OS ONE
Traps (notify_type = traps)	Yes (v2c and v3)	Yes (v2c and v3), optional
Informs (notify_type = informs)	v2c: rejected; v3: yes	v2c: yes; v3: yes
Engine ID	v3 informs only	v3 informs only
Maximum hosts	6 per version	18 total
SNMPv2c informs	Not supported	Supported (no engine ID)

API Reference

Proxy Path — GoSNMP Service (Port 8092)

Purpose	Method	Endpoint
Register a proxy subscriber	POST	/v1/snmp/subscribers
List proxy subscribers	GET	/v1/snmp/subscribers
Delete a proxy subscriber	DELETE	/v1/snmp/subscriber/{host}
Check service health	GET	/v1/snmp/health
List forwarding executions	GET	/v1/snmp/executions

Direct Path — GoInventory Service (Port 8080)

Purpose	Method	Endpoint
Create / list / delete a direct host	POST / GET / DELETE	/v1/inventory/snmp/host
Create / list / delete a view	POST / GET / DELETE	/v1/inventory/snmp/view
Create / list / delete a user	POST / GET / DELETE	/v1/inventory/snmp/user
Create / list / delete a community	POST / GET / DELETE	/v1/inventory/snmp/community
List / update a group	GET / PUT	/v1/inventory/snmp/group
Create / list / delete a use-vrf	POST / GET / DELETE	/v1/inventory/snmp/use-vrf

Subscriber Fields (Proxy Path, Port 8092)

Field	Type	Constraints	Description
host	string	required	IPv4/IPv6 address or hostname of the NMS.
id	integer	read-only	Auto-assigned identifier.
type	string	required; v2c or v3	Subscriber protocol version.
v2community	string	—	SNMPv2c community string.
v3user	string	—	SNMPv3 username.
v3authprotocol	string	md5 or sha1	Authentication protocol.
v3authpassword	string	—	Authentication passphrase.
v3privprotocol	string	aes or des	Privacy protocol.
v3privpassword	string	—	Privacy passphrase.

Host Fields (Direct Path, Port 8080)

Field	Type	Values / Constraints	Description
host_ip	string	IPv4/IPv6/FQDN	External manager address.
community	string	—	SNMP community name (v2c).
user	string	—	SNMP user name (v3).

Field	Type	Values / Constraints	Description
version	string	v2c or v3	SNMP version.
engineid	string	hex, colon-separated	Remote engine ID (for informs).
notify_type	string	traps or informs	Notification type.
udp_port	string	default 162	UDP destination port.

Configure Model A: XCO Proxy (Default)

The proxy path is the fastest way to get notifications flowing. XCO configures your switches automatically; you only need to tell XCO where your NMS is.

About This Task

When you register a switch in XCO 4.1.0, XCO pushes a default SNMP configuration to the switch so it sends its notifications to XCO. This includes a default SNMP view, group, user, and host, and works on both SLX-OS and OS ONE. When you unregister the device, XCO removes this configuration for you.

Note

Default views and groups are created during registration starting in XCO 4.1.0. On the proxy path, the auto-created host sends traps only. If you need informs, use the Direct path (Model B).

What NMS Will Receive

XCO re-encodes each notification before forwarding it. Expect the following:

Constraint	Behavior
Maximum subscribers	4 system-wide. Registering a fifth returns an error.
Delivery guarantee	Best-effort — no acknowledgment and no retry.
Informs	Not supported; the proxy always sends traps to subscribers.
Per-switch filtering	Not available; every subscriber receives notifications from every switch.
Subscriber SNMP versions	v2c and v3.
Duplicate subscribers	Each (type, host) pair must be unique; re-registering the same host returns a conflict error.
Forwarding order	Sequential — a slow or unresponsive subscriber delays delivery to the ones after it.

Proxy Path Limits

Constraint	Behavior
Maximum subscribers	4 system-wide. Registering a fifth returns an error.
Delivery guarantee	Best-effort — no acknowledgment and no retry.

Constraint	Behavior
Informs	Not supported; the proxy always sends traps to subscribers.
Per-switch filtering	Not available; every subscriber receives notifications from every switch.
Subscriber SNMP versions	v2c and v3.
Duplicate subscribers	Each (type, host) pair must be unique; re-registering the same host returns a conflict error.
Forwarding order	Sequential — a slow or unresponsive subscriber delays delivery to the ones after it.

Follow this procedure to configure proxy path to get the notification.

Procedure

1. Register your NMS as a subscriber.

   Send a POST request to the proxy service on port 8092. SNMPv3 is recommended because it authenticates and encrypts your notifications.

   SNMPv3 subscriber (recommended)

   ```
   curl -X POST https://<XCO_IP>:8092/v1/snmp/subscribers \
     -H "Content-Type: application/json" \
     -H "Authorization: Bearer <token>" \
     -d '{
       "type": "v3", "host": "10.0.0.50", "v3user": "nms_user",
       "v3authprotocol": "sha1", "v3authpassword": "MyAuthP@ss123",
       "v3privprotocol": "aes",  "v3privpassword": "MyPrivP@ss456"
     }
   ```

   SNMPv2c subscriber (simpler, less secure)

   ```
   curl -X POST https://<XCO_IP>:8092/v1/snmp/subscribers \
     -d '{"type": "v2c", "host": "10.0.0.50", "v2community": "public"}'
   ```

2. Verify the registration.

   ```
   curl -X GET https://<XCO_IP>:8092/v1/snmp/subscribers \
     -H "Authorization: Bearer <token>"
   ```

   

   Tip
   Once registered, your NMS receives traps from every switch that XCO manages. The proxy path does not filter by switch — all subscribers receive notifications from all managed switches.

3. Day-2 operations

   a. Remove a subscriber.

   ```
   curl -X DELETE https://<XCO_IP>:8092/v1/snmp/subscriber/10.0.0.50 \
     -H "Authorization: Bearer <token>"
   ```

   b. Check audit logs of forwarding activity.

   ```
   curl -X GET https://<XCO_IP>:8092/v1/snmp/executions \
     -H "Authorization: Bearer <token>"
   ```

Configure Model B: Direct to External Manager

You can configure Model B.

About This Task

On the Direct path there is no automatic setup. You configure each switch through the GoInventory service on port 8080 to send notifications straight to your NMS. The exact steps depend on the switch platform.

Both platforms follow the same required order: create a View, then a Group (with notify access), then a User or Community, and finally a Host. Each element must exist before the next one can be created. If your NMS is only reachable over a non-default VRF, configure the VRF as well.

Note
XCO's automatic proxy setup already creates efav3View, efav3Group, and efav3User on each switch. You can reuse these objects for direct hosts (simplest), or create your own for full isolation.

Follow this procedure to configure direct path to get the notification.

Procedure

1. On SLX-OS devices.
   a. (Option A) Reuse XCO's existing objects (simplest).

      The existing `efav3Group` already has notify access, and `efav3View` already covers the 1.3.6.1 subtree. Skip straight to creating the host (Step 4 of Option B). If your NMS is only reachable via a non-default VRF, also complete the VRF step.
   b. (Option B) Create your own SNMP objects (full isolation) — SNMPv3.
      i. Create an SNMP view.

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/view \
  -d '{"device_ips": ["10.1.1.1"],
       "view": {"name": "nmsView", "oid": "1.3.6.1", "mib_access": "included"}}'
```

      ii. Update the SNMP group to enable notify access (required for any host, traps or informs)

```
curl -X PUT https://<XCO_IP>:8080/v1/inventory/snmp/group \
  -d '{"device_ips": ["10.1.1.1"],
       "group": {"name": "nmsGroup", "version": "v3",
                 "auth_level": "auth_with_privacy",
                 "read_view": "nmsView", "write_view": "nmsView",
                 "notify_view": "nmsView"}}'
```

Note
The group's
notify_view
must be set. Without it, the switch sends no notifications no matter how the host is configured. Valid **auth_level** values are auth_disabled (no auth), auth_with_no_privacy (auth only), and auth_with_privacy (auth and encryption). This field does not apply to v2c groups.

iii. Create the SNMPv3 user

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/user \
  -d '{"device_ips": ["10.1.1.1"],
       "user": {"name": "nmsUser", "group_name": "nmsGroup", "version": "v3",
                "auth_protocol": "sha", "auth_password": "<password>",
                "priv_protocol": "aes128", "priv_password": "<password>"}}
```

iv. Create the host. For traps (fire-and-forget)

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/host \
  -d '{"device_ips": ["10.1.1.1"],
       "host": {"host_ip": "10.0.0.50", "user": "nmsUser",
                "version": "v3", "notify_type": "traps", "udp_port": "162"}}'
```

For informs (reliable, with acknowledgment and retry — SNMPv3 only, requires your NMS's engine ID):

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/host \
  -d '{"device_ips": ["10.1.1.1"],
       "host": {"host_ip": "10.0.0.50", "user": "nmsUser", "version": "v3",
                "notify_type": "informs",
                "engineid": "80:00:13:70:01:c0:a8:01:64", "udp_port": "162"}}'
```

Note
The engine ID for informs must be your NMS's SNMP engine ID — not the switch's. The switch uses it to encrypt the inform request for the manager. Find it in your NMS configuration, or generate one per RFC 3411. Use colon-separated hex octets.

c. (Option C) Use SNMPv2c instead of v3 (SLX-OS).
For v2c direct delivery, replace Steps 2–4 above with a community and a host. Note that SNMPv2c on SLX-OS supports traps only — not informs.

Create the SNMP community (this creates an implicit group)

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/community \
  -d '{"device_ips": ["10.1.1.1"],
       "community": {"name": "nms_community", "group_name": "nmsGroup",
                     "enable_notify_access": true, "view": "nmsView"}}'
```

Create the host that uses the community

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/host \
  -d '{"device_ips": ["10.1.1.1"],
       "host": {"host_ip": "10.0.0.50", "community": "nms_community",
                "version": "v2c", "notify_type": "traps", "udp_port": "162"}}'
```

d. (Optional) Configure a VRF (SLX-OS)
Complete this step only if your NMS is reachable solely through a non-default VRF:

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/use-vrf \
  -d '{"device_ips": ["10.1.1.1"], "use-vrf": {"name": "vrf1"}}'
```

2. On OS ONE devices.
OS ONE follows the same required order — `View → Group → User/Community → Host` — using the same GoInventory endpoints on port 8080. The differences from SLX-OS are:

- **Host capacity:** up to 18 hosts total (v2c and v3 combined), versus 6 per version on SLX-OS.
- **v2c informs:** supported on OS ONE (no engine ID required); on SLX-OS, informs require SNMPv3.
- **User protocols:** OS ONE also accepts noauth (no authentication) and nopriv (no encryption); SLX-OS requires md5/sha for authentication and des/aes128 for privacy.

The following example configures an SNMPv3 inform host on OS ONE. Use your NMS's engine ID for informs, just as on SLX-OS:

```
curl -X POST https://<XCO_IP>:8080/v1/inventory/snmp/host \
  -d '{"device_ips": ["10.2.2.1"],
       "host": {"host_ip": "10.0.0.50", "user": "nmsUser", "version": "v3",
                "notify_type": "informs",
                "engineid": "80:00:13:70:01:c0:a8:01:64", "udp_port": "162"}}'
```

Configure Model C: Hybrid (Both Paths Simultaneously)

You can configure hybrid model.

About This Task

Use the hybrid model when you want central aggregation through the XCO proxy — for dashboards and event correlation — together with reliable, direct delivery to a compliance NMS. The two paths run independently: the switch sends notifications to XCO and to the direct host at the same time.

Follow this procedure to set up hybrid model. Complete both Model A and Model B. No extra coordination between the paths is needed.

Procedure

1. Configure Model A.

   See Configure Model A: XCO Proxy (Default) on page 881.
2. Configure Model B.

   See Configure Model B: Direct to External Manager on page 883.

XCO SNMP Support for Extreme OS ONE

The SNMP service enables communication between network devices and a central manager. Here's how it works with OS ONE:

- XCO acts as the SNMP Manager, receiving alert messages (traps) from OS ONE devices.
- When an OS ONE device is registered with XCO, it's automatically configured to send SNMP v3 traps to XCO.
- XCO can forward received traps from OS ONE devices to an external receiver, acting as an SNMP proxy for SNMP v2 and v3 traps.
- To set up trap reception, XCO requires specific credentials (v3 username, authentication key, and privacy key, or a v2c community string).
- During updates, XCO verifies its registration status with OS ONE devices and updates SNMP configurations accordingly. If a device is unregistered, it stops sending traps to XCO.

Configure SNMP

You can configure SNMP on OS ONE.

About This Task

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 OS ONE devices and agents, receiving traps from all devices in its inventory. Once you register a OS ONE device with XCO, XCO automatically configures the OS ONE device to send v3 traps to XCO.

XCO acts as an SNMP proxy for all SNMP v2 and v3 traps received from OS ONE devices, forwarding them to an external trap receiver if one is present.

- 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 OneOS 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.

For information on command syntax and parameter description, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Follow this procedure to configure SNMP.

Procedure

1. Configure SNMP View
2. Configure SNMP Community on page 886
3. Configure SNMP Group
4. Configure SNMP User on page 889
5. Configure SNMP Host on page 894
6. Configure SNMP USE VRF on page 899

Configure SNMP Community

You can configure SNMP Community on OS ONE.

About This Task
Follow this procedure to configure the SNMP community.

Procedure

1. Create an SNMP community.
   a. Run the **`efa inventory device snmp view create`** command on SLX devices.

```
efa inventory device snmp view create --ip 10.64.160.30 --name view1 --mib-tree
1.3.6.1 --mib-tree-access included
+--------------+-------+----------+-----------------+---------+---------+
```

```
|  IP Address  | Name  | MIB-Tree | MIB-Tree-Access | Status  | Reason |
+--------------+-------+----------+-----------------+---------+--------+
| 10.64.160.30 | view1 | 1.3.6.1  | included        | Success |        |
+--------------+-------+----------+-----------------+---------+--------+
Snmp view details
--- Time Elapsed: 10.775052081s ---
```

b. Run the **show running-config snmp-server** command on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf mgmt-
vrf


snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify efav3View
```

2. List SNMP community.
   a. Run the **efa inventory device snmp view list** command.

```
efa inventory device snmp view list --ip 10.64.160.30
+--------------+-----------+----------+-----------------+-------------+
|  IP Address  |   Name    | MIB-Tree | MIB-Tree_Access |  AppState   |
+--------------+-----------+----------+-----------------+-------------+
| 10.64.160.30 | efav3View | 1.3.6.1  | included        | cfg-in-sync |
+              +-----------+----------+-----------------+-------------+
|              | view1     | 1.3.6.1  | included        | cfg-in-sync |
+--------------+-----------+----------+-----------------+-------------+
Snmp view details
--- Time Elapsed: 17.568431ms ---
```

3. Delete SNMP community.
   a. Run the **efa inventory device snmp community delete** command.

```
efa inventory device snmp community delete --ip 10.64.160.30 --name comm1
+--------------+-------------------------------+---------+--------+
|  IP Address  |       Community Name          | Status  | Reason |
+--------------+-------------------------------+---------+--------+
| 10.64.160.30 | $9$dJvAw/ns+vCSE21/mR8V7A== | Success |        |
+--------------+-------------------------------+---------+--------+
Snmp community details
--- Time Elapsed: 11.286033594s ---
```

   b. Run the **show running-config snmp-server** command on SLX devices.

```
SLX-30# show running-config snmp-
server
```

```
snmp-server sys-descr "Extreme 8720
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf mgmt-
vrf


snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v3 priv
```

c. Run the **`efa inventory device snmp community delete`** command on OS ONE devices.

```
efa inventory device snmp community delete --ip 10.64.184.88 --name comm1
+--------------+----------------+---------+--------+
|  IP Address  | Community Name | Status  | Reason |
+--------------+----------------+---------+--------+
| 10.64.184.88 | comm1          | Success |        |
+--------------+----------------+---------+--------+
Snmp community details
--- Time Elapsed: 10.46719085s ---
```

d. Run the **`show running-config system snmp-server`** command on SLX devices.

```
show running-config system snmp-
server


system


  snmp-server
DEFAULT


    vrf mgmt-
vrf



enable


    !
```

```
    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b

  !


!
```

Configure SNMP User

You can configure SNMP user on OS ONE.

About This Task

Follow this procedure to configure SNMP user.

Procedure

1. Create an SNMP user.
   a. Run the **`efa inventory device snmp user create`** command.

```
efa inventory device snmp user create --ip 10.64.160.30 --name user1 --group
group2 --auth-protocol sha --auth-pass Password@1 --priv-protocol AES128 --priv-
pass Password@1 --view view1
+--------------+-------+--------+-----------+------------+-------------
+-------+------------+-------------------------------+------------
+-------------------------------+---------+--------+
|  IP Address  | User  | Group
| Read view | Write view | Notify view | View  | Auth proto |       Auth
passphrase        | Priv proto |       Priv passphrase        | Status  | Reason |
+--------------+-------+--------+-----------+------------+-------------
+-------+------------+-------------------------------+------------
+-------------------------------+---------+--------+
| 10.64.160.30 | user1 | group2 |
|              |              | view1 | sha        | WKRubRPg45DNTqLQKAmsV8DMxFw=
| AES128     | WKRubRPg45DNTqLQKAmsV8DMxFw= | Success |        |
+--------------+-------+--------+-----------+------------+-------------
+-------+------------+-------------------------------+------------
+-------------------------------+---------+--------+
Snmp user details
    --- Time Elapsed: 11.472501294s ---
```

   b. Run the **`show running-config snmp-server`** command on SLX devices.

```
SLX-30# show running-config snmp-
server


     snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf mgmt-
vrf


snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
```

```
encrypted
snmp-server user user1 groupname group2 auth sha auth-password
WKRubRPg45DNTqLQKAmsV8DMxFw= priv AES128 priv-password WKRubRPg45DNTqLQKAmsV8DMxFw=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c notify
view1


snmp-server group group1 v3
priv


snmp-server group group2 v3 priv
```

c. Run the **`efa inventory device snmp user create`** command on OS ONE devices.

```
efa inventory device snmp user create --ip 10.64.184.88 --name user1 --auth-
protocol sha --auth-pass Password@1 --priv-protocol aes --priv-pass Password@1
+--------------+-------+-------+-----------+------------+-------------+------
+------------+-----------------------------------------------+-------------
+---------------------------------------------+---------+--------+
|  IP Address  | User  | Group | Read view | Write
view | Notify view | View | Auth proto |              Auth passphrase
| Priv proto |              Priv passphrase               | Status  | Reason |
+--------------+-------+-------+-----------+------------+-------------+------
+------------+-----------------------------------------------+-------------
+---------------------------------------------+---------+--------+
| 10.64.184.88 | user1 |       |           |
|              |      | sha        | 0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c
| aes        | 0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c | Success |        |
+--------------+-------+-------+-----------+------------+-------------+------
+------------+-----------------------------------------------+-------------
+---------------------------------------------+---------+--------+
Snmp user details
--- Time Elapsed: 11.224769071s ---
```

d. Run the **`show running-config system snmp-server`** command on SLX devices.

```
vtos88# show running-config system snmp-
server


system


  snmp-server
DEFAULT


    vrf mgmt-
vrf
```

```
enable


    !


    community-hashed
zfrOmT+rK3


    user user1 auth sha auth-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c priv aes priv-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c

    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b

  !


!
```

2. List SNMP user.
   a. Run the **`efa inventory device snmp user list`** command.

```
efa inventory device snmp user list --ip 10.64.160.30
+--------------+-------+--------+-----------+------------+-------------
+-------+------------+-------------------------------+------------
+-------------------------------+-------------+
|  IP Address  | User  | Group
| Read view | Write view | Notify view | View  | Auth proto |
Auth passphrase        | Priv proto |       Priv passphrase        |  AppState   |
+--------------+-------+--------+-----------+------------+-------------
+-------+------------+-------------------------------+------------
+-------------------------------+-------------+
| 10.64.160.30 | user1 | group2 |
|            |             | view1 | sha        | WKRubRPg45DNTqLQKAmsV8DMxFw=
| AES128     | WKRubRPg45DNTqLQKAmsV8DMxFw= | cfg-in-sync |
+--------------+-------+--------+-----------+------------+-------------
+-------+------------+-------------------------------+------------
+-------------------------------+-------------+
Snmp user details
--- Time Elapsed: 47.932936ms ---
```

   b. Run the **`efa inventory device snmp user list`** command on OS ONE devices.

```
efa inventory device snmp user list --ip 10.64.184.88
+--------------+-------+-------+-----------+------------+-------------+------
+------------+-----------------------------------------------+------------
+-------------------------------------------+-------------+
|  IP Address  | User  | Group | Read view | Write
view | Notify view | View | Auth proto |              Auth passphrase
| Priv proto |              Priv passphrase               |  AppState   |
+--------------+-------+-------+-----------+------------+-------------+------
+------------+-----------------------------------------------+------------
+-------------------------------------------+-------------+
| 10.64.184.88 | user1 |       |           |
|              |      | sha        | 0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c
| aes        | 0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c | cfg-in-sync |
+--------------+-------+-------+-----------+------------+-------------+------
+------------+-----------------------------------------------+------------
+-------------------------------------------+-------------+
Snmp user details
--- Time Elapsed: 13.671279ms ---
```

3. Delete SNMP view.
   a. Run the **efa inventory device snmp view delete** command.

```
efa inventory device snmp view delete --ip 10.64.160.30 --name view1
+--------------+-------+----------+---------+--------+
|  IP Address  | Name  | MIB-Tree | Status  | Reason |
+--------------+-------+----------+---------+--------+
| 10.64.160.30 | view1 | 1.3.6.1  | Success |        |
+--------------+-------+----------+---------+--------+
Snmp view details
--- Time Elapsed: 10.78019127s ---
```

   b. Run the **show running-config snmp-server** command on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf mgmt-
vrf


snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v3 priv
```

4. Delete SNMP user.
   a. Run the **efa inventory device snmp user delete** command on SLX devices.

```
efa inventory device snmp user delete --ip 10.64.160.30 --name user1
+--------------+-----------+---------+--------+
|  IP Address  | User Name | Status  | Reason |
+--------------+-----------+---------+--------+
| 10.64.160.30 | user1     | Success |        |
+--------------+-----------+---------+--------+
Snmp user details
--- Time Elapsed: 11.348887577s ---
```

   b. Complete the following configuration on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5
```

```
snmp-server use-vrf mgmt-
vrf


snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c notify
view1


snmp-server group group1 v3 priv
```

c. Run the **`efa inventory device snmp user delete`** command on OS ONE devices.

```
efa inventory device snmp user delete --ip 10.64.184.88 --name user1
+--------------+-----------+---------+--------+
|  IP Address  | User Name | Status  | Reason |
+--------------+-----------+---------+--------+
| 10.64.184.88 | user1     | Success |        |
+--------------+-----------+---------+--------+
Snmp user details
--- Time Elapsed: 10.703787345s ---
```

d. Complete the following configuration on OS ONE devices.

```
vtos88# show running-config system snmp-
server


system


  snmp-server
DEFAULT


    vrf mgmt-
vrf



enable


    !


    community-hashed
zfrOmT+rK3
```

```
    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b

  !


!
```

Configure SNMP Host

You can configure SNMP host on OS ONE.

About This Task
Follow this procedure to configure SNMP host.

Procedure

1. Create an SNMP host.
   a. Run the **`efa inventory device snmp host create`** command.

```
efa inventory device snmp host create --ip 10.64.160.30 --host-ip 1.1.1.1 --
community comm1 --version v2c --notify-type traps
+--------------+---------+------+------------------------------+-------------
+----------+----------------------+----------+----------+----------+---------
+--------+
|  IP Address  | Host IP | User |          Community
| Notify Type | EngineID |   Source Interface    |
Vrf    | UDP port | Severity | Status  | Reason |
+--------------+---------+------+------------------------------+-------------
+----------+----------------------+----------+----------+----------+---------
+--------+
| 10.64.160.30 | 1.1.1.1 |      | $9$dJvAw/ns+vCSE21/mR8V7A==
| traps       |          | management chassis-ip | mgmt-vrf
| 162      | None     | Success |        |
+--------------+---------+------+------------------------------+-------------
+----------+----------------------+----------+----------+----------+---------
+--------+
Snmp host details
--- Time Elapsed: 10.753050364s ---
```

   b. Complete the following configuration on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf default-
vrf


snmp-server use-vrf mgmt-
vrf
```

```
snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server host 1.1.1.1 $9$dJvAw/ns+vCSE21/
mR8V7A==

 version
2c


 source-interface management chassis-
ip


!


snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server user user1 groupname group2 auth sha auth-password
WKRubRPg45DNTqLQKAmsV8DMxFw= priv AES128 priv-password WKRubRPg45DNTqLQKAmsV8DMxFw=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c notify
view1


snmp-server group group1 v3
priv


snmp-server group group2 v3 priv
```

c. Run the **`efa inventory device snmp host create`** command on OS ONE devices.

```
efa inventory device snmp host create --ip 10.64.184.88 --host-ip 1.1.1.1 --
community comm1 --version v2c
+--------------+---------+------+-----------+-------------+----------
+------------------+-----+----------+----------+---------+--------+
|  IP Address  | Host IP | User | Community | Notify
Type | EngineID | Source Interface | Vrf | UDP port | Severity | Status  | Reason |
+--------------+---------+------+-----------+-------------+----------
+------------------+-----+----------+----------+---------+--------+
| 10.64.184.88 | 1.1.1.1 |      | comm1     |
|          |                  |     | 162      |          | Success |        |
+--------------+---------+------+-----------+-------------+----------
+------------------+-----+----------+----------+---------+--------+
Snmp host details
--- Time Elapsed: 10.297739324s ---
```

d. Complete the following configuration on OS ONE devices.

```
vtos88# show running-config system snmp-
server
```

```
system


  snmp-server
DEFAULT


    vrf mgmt-
vrf



enable


      host 1.1.1.1
zfrOmT+rK3


        version
2c


        udp-port
162


      !


    !


    vrf default-
vrf



enable


    !


    community-hashed
zfrOmT+rK3


    user user1 auth sha auth-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c

0ccc45c priv aes priv-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c


    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc72

1cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5
```

```
de2b


  !


!
```

2. List SNMP host.
   a. Run the **efa inventory device snmp host list** command.

```
efa inventory device snmp host list --ip 10.64.184.88
+--------------+---------+------+-----------+-------------+-----------------
+------------------+-----+----------+----------+-------------+
|  IP Address  | Host IP | User | Community | Notify Type
| Remote EngineID | Source Interface | Vrf | UDP port | Severity |  AppState   |
+--------------+---------+------+-----------+-------------+-----------------
+------------------+-----+----------+----------+-------------+
| 10.64.184.88 | 1.1.1.1 |      | comm1     |
|                 |                  |     | 162      |          | cfg-in-sync |
+--------------+---------+------+-----------+-------------+-----------------
+------------------+-----+----------+----------+-------------+
Snmp host details
--- Time Elapsed: 18.649658ms ---
```

   b. Run the **efa inventory device snmp host list** command on OS ONE devices.

```
efa inventory device snmp host list --ip 10.64.160.30
+--------------+---------+------+-------------------------------+-------------
+-----------------+-----------------------+----------+----------+----------
+-------------+
|  IP Address  | Host IP | User |          Community
| Notify Type | Remote EngineID |   Source Interface
|   Vrf    | UDP port | Severity |  AppState   |
+--------------+---------+------+-------------------------------+-------------
+-----------------+-----------------------+----------+----------+----------
+-------------+
| 10.64.160.30 | 1.1.1.1 |      | $9$dJvAw/ns+vCSE21/mR8V7A==
| traps       |                 | management chassis-ip |
mgmt-vrf | 162      | None     | cfg-in-sync |
+--------------+---------+------+-------------------------------+-------------
+-----------------+-----------------------+----------+----------+----------
+-------------+
Snmp host details
--- Time Elapsed: 17.261472ms ---
```

3. Delete SNMP host.
   a. Run the **efa inventory device snmp host delete** command.

```
efa inventory device snmp host delete --ip 10.64.160.30 --host-ip 1.1.1.1 --
community comm1
+--------------+---------+---------+--------+
|  IP Address  | Host IP | Status  | Reason |
+--------------+---------+---------+--------+
| 10.64.160.30 | 1.1.1.1 | Success |        |
+--------------+---------+---------+--------+
Snmp host details
--- Time Elapsed: 10.765275894s ---
```

   b. Complete the following configuration on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
```

```
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf default-
vrf


snmp-server use-vrf mgmt-
vrf


snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server user user1 groupname group2 auth sha auth-password
WKRubRPg45DNTqLQKAmsV8DMxFw= priv AES128 priv-password WKRubRPg45DNTqLQKAmsV8DMxFw=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c notify
view1


snmp-server group group1 v3
priv


snmp-server group group2 v3 priv
```

c. Run the **`efa inventory device snmp host delete`** command on OS ONE devices.

```
efa inventory device snmp host delete --ip 10.64.184.88 --host-ip 1.1.1.1 --
community comm1
+--------------+---------+---------+--------+
|  IP Address  | Host IP | Status  | Reason |
+--------------+---------+---------+--------+
| 10.64.184.88 | 1.1.1.1 | Success |        |
+--------------+---------+---------+--------+
Snmp host details
--- Time Elapsed: 10.761429248s ---
```

d. Complete the following configuration on OS ONE device.

```
vtos88# show running-config system snmp-
server


system
```

```
  snmp-server
DEFAULT


    vrf mgmt-
vrf



enable


    !


    vrf default-
vrf



enable


    !


    community-hashed
zfrOmT+rK3


    user user1 auth sha auth-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c priv aes priv-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c

    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b

  !


!
```

Configure SNMP USE VRF

You can configure SNMP USE VRF on OS ONE.

About This Task

Follow this procedure to configure SNMP USE VRF.

Procedure

1. Create an SNMP USE VRF.
   a. Run the **`efa inventory device snmp use-vrf create`** command.

```
efa inventory device snmp use-vrf create --ip 10.64.160.30 --name default-vrf
+--------------+-------------+----------+---------+--------+
|  IP Address  |   Use VRF   | Shutdown | Status  | Reason |
+--------------+-------------+----------+---------+--------+
| 10.64.160.30 | default-vrf | false    | Success |        |
+--------------+-------------+----------+---------+--------+
```

```
Snmp use-vrf details
            --- Time Elapsed: 10.773595273s ---
```

b. Complete the following configuration on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf default-
vrf


snmp-server use-vrf mgmt-
vrf


snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server user user1 groupname group2 auth sha auth-password
WKRubRPg45DNTqLQKAmsV8DMxFw= priv AES128 priv-password WKRubRPg45DNTqLQKAmsV8DMxFw=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c notify
view1


snmp-server group group1 v3
priv


snmp-server group group2 v3 priv
```

c. Run the **`efa inventory device snmp use-vrf create`** command on OS ONE devices.

```
efa inventory device snmp use-vrf create --ip 10.64.184.88 --name default-vrf
+--------------+-------------+----------+---------+--------+
|  IP Address  |   Use VRF   | Shutdown | Status  | Reason |
+--------------+-------------+----------+---------+--------+
| 10.64.184.88 | default-vrf | false    | Success |        |
+--------------+-------------+----------+---------+--------+
```

```
Snmp use-vrf details
--- Time Elapsed: 10.539498207s ---
```

d. Complete the following configuration on OS devices.

```
vtos88# show running-config system snmp-
server



system


      snmp-server
DEFAULT


        vrf mgmt-
vrf



enable


        !


        vrf default-
vrf



enable


        !


        community-hashed
zfrOmT+rK3


 user user1 auth sha auth-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c priv aes priv-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c

 user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b

  !


!
```

2. List SNMP USE VRF.

a. Run the **`efa inventory device snmp use-vrf list`** command.

```
efa inventory device snmp use-vrf list --ip 10.64.160.30
+--------------+-------------+----------+-------------+
|  IP Address  |   Use VRF   | Shutdown |  AppState   |
+--------------+-------------+----------+-------------+
| 10.64.160.30 | default-vrf | false    | cfg-in-sync |
+--------------+-------------+----------+-------------+
```

```
Snmp use-vrf details
--- Time Elapsed: 15.831853ms ---
```

b. Run the **`efa inventory device snmp use-vrf list`** command on OS ONE devices.

```
efa inventory device snmp use-vrf list --ip 10.64.184.88
+--------------+-------------+----------+-------------+
|  IP Address  |   Use VRF   | Shutdown |  AppState   |
+--------------+-------------+----------+-------------+
| 10.64.184.88 | default-vrf | false    | cfg-in-sync |
+--------------+-------------+----------+-------------+
Snmp use-vrf details
--- Time Elapsed: 19.874455ms ---
```

3. Delete SNMP USE VRF.

a. Run the **`efa inventory device snmp use-vrf delete`** command.

```
efa inventory device snmp use-vrf delete --ip 10.64.160.30 --name default-vrf
+--------------+-------------+---------+--------+
|  IP Address  |  User Name  | Status  | Reason |
+--------------+-------------+---------+--------+
| 10.64.160.30 | default-vrf | Success |        |
+--------------+-------------+---------+--------+
Snmp user details
--- Time Elapsed: 10.866799854s ---
```

b. Complete the following configuration on SLX devices.

```
SLX-30# show running-config snmp-
server


snmp-server sys-descr "Extreme 8720-32C
Switch"

snmp-server engineID local
80:0:6:34:b2:9b:48:0:10:5:f0:7f:d5


snmp-server use-vrf mgmt-
vrf


snmp-server community $9$dJvAw/ns+vCSE21/mR8V7A== groupname
group1

snmp-server user efav3User groupname efav3Group auth sha auth-password
LlgMJxy1dZfAfPydPDdFTNNtkpw= priv AES128 priv-password 53IJSJKiK/A6UFLVd9+ZkkclIJs=
encrypted
snmp-server user user1 groupname group2 auth sha auth-password
WKRubRPg45DNTqLQKAmsV8DMxFw= priv AES128 priv-password WKRubRPg45DNTqLQKAmsV8DMxFw=
encrypted
snmp-server view efav3View 1.3.6.1
included


snmp-server view view1 1.3.6.1
included


snmp-server group efav3Group v3 notify
efav3View


snmp-server group group1 v2c notify
view1
```

```
snmp-server group group1 v3
priv

snmp-server group group2 v3 priv
```

c. Run the **efa inventory device snmp use-vrf delete** command on OS ONE devices.

```
efa inventory device snmp use-vrf delete --ip 10.64.184.88 --name default-vrf
+--------------+-------------+---------+--------+
|  IP Address  |  User Name  | Status  | Reason |
+--------------+-------------+---------+--------+
| 10.64.184.88 | default-vrf | Success |        |
+--------------+-------------+---------+--------+
Snmp user details
--- Time Elapsed: 10.358329034s ---
```

d. Complete the following configuration on OS ONE devices.

```
vtos88# show running-config system snmp-
server

system

  snmp-server
DEFAULT

    vrf mgmt-
vrf

enable

    !

    community-hashed
zfrOmT+rK3

    user user1 auth sha auth-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c priv aes priv-password-hashed
0x58a46e6d13e0e390cd4ea2d02809ac57c0ccc45c

    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b

  !

!
```

Host Operating System (Host OS) Event Logging Configuration

Use this topic to learn about the logging security events on the Host Operating System (OS). Use Auditd, a Linux Auditing System component, to record and track system-level events. This provides detailed logging of security-relevant information, enabling administrators to monitor and analyze potential security incidents and ensure policy compliance.

The security event logging ensures the following benefit:

- Automated Auditd Installation: Install Auditd automatically during XCO installation for seamless setup and consistent system auditing.
- Default Audit Rules: Include a hardening script default audit.rules file with fundamental and critical logging categories to ensure essential security events are logged. A hardening script must be executed for XCO TPVM deployment.
- Log Forwarding: Configure settings via XCO (either using CLI or API) to forward selected audit log categories and audit log record types to external syslog or webhook servers for centralized analysis from multiple sources.
- Log Storage Management: Configure storage limits and implement log rotation to prevent logs from exceeding available storage space, ensuring optimal system performance.

Auditd Installation and Default Audit Rules

During the standard installation process, the Auditd service is automatically installed to enable auditing capabilities. For TPVM deployments, TPVM installs auditd service and its required packages. For server deployments, XCO installs auditd along with it's dependencies.

XCO ensures that Auditd is installed, activated, and configured to forward audit logs.

A TPVM hardening script is provided by XCO to configure default audit rules, including CIS-CAT industry standard audit rules. To add these rules, you must manually run this script after XCO installation.

Post-Installation Audit Rule Management

Post installation of XCO and Auditd, you can modify, extend, and manage audit rules to meet their organization's specific requirements and compliance standards.

Note

- XCO is not responsible for managing audit rules post-installation.
- Default rules are part of XCO hardening script.
- Auditd generates default logs, regardless of whether custom or default user audit rules are in place.

Audit Log Delivery Process

The audit logs are delivered in the following order:

1. Auditd service generates audit logs using XCO hardening script default audit rule file.
2. Logs are forwarded to the rsyslog service on the host OS.
3. The host OS syslog client sends audit logs to the RASLog Service.
4. RASLog Service forwards audit logs to the Notification service.
5. XCO Notification filters all the audit logs based on subscriber filters and sends them to external webhook servers or syslog servers via secure or insecure medium.

Notification CLI Changes for Audit Log Filtering

The following examples shows how to enable Syslog subscribers for various audit log filter types:

- Enable a Syslog subscriber for HOST_EVENT filtering, tailored to Auditd events. This configuration targets specific audit rule group key "logins" and record types (USER_LOGIN and USER_END), transmitting data over an insecure channel.

```
(efa:extreme)extreme@tpvm:~$ efa notification subscribers add-syslog-relp --address
127.0.0.1:1601 --insecure --filter HOST_EVENTS --host-event auditd --host-auditd-key
logins --minimum-severity warning --host-auditd-record-type USER_LOGIN,USER_END
Successfully registered subscriber.

+-----------
+-------------------------------------------------------------------------------------+
| attribute |
value                                                                                |
+-----------
+-------------------------------------------------------------------------------------+
| id        |
13                                                                                   |
+-----------
+-------------------------------------------------------------------------------------+
| handler   |
relp                                                                                 |
+-----------
+-------------------------------------------------------------------------------------+
| endpoint
| 127.0.0.1:1601                                                                       |
+-----------
+-------------------------------------------------------------------------------------+
| config
| {"cacert":"","conn-timeout":10,"device-event":[],"filters":["HOST_EVENTS"],"host|
|           |
-auditd-key":["logins"],"host-auditd-record-type":["USER_LOGIN","USER_END"],     |
|           |
"host-event":["auditd"],"insecure":true,"minimum-severity":"warning              |
|           |
","rfc5424":false}                                                                   |
+-----------
+-------------------------------------------------------------------------------------+
Notification Subscriber ID=13
--- Time Elapsed: 2.047374836s ---

### Values for filters:

* "DEVICE_EVENTS" - Syslog/Raslog generated from devices
* "HOST_EVENTS " - Events generated from the HOST OS
* "APP_EVENTS" - Task or user events generated from the application
```

```
* "APP_ALERTS" - Fault alerts generated from the application
* "APP_ALARMS" - Fault alarms generated from the application

### Values for host-event:

* "auditd" - Audit log messages

### host-auditd-key (Optional)
The host-auditd-key
is a custom key representing an audit rule group. This can either belong to the
default audit rules provided or to any new audit rule group introduced by the user.
Example:
logins, time-change
logins: Represents the audit rule group for login-related events.

time-change: Represents
the audit rule group for time synchronization or system time changes.

Users can define
or reference such keys to tailor the audit configuration to meet organizational needs

### host-auditd-record-type (Optional)

The host-auditd-record-type
is a standard message/record type representing an audit event.
Example:
USER_LOGIN        User login events, successful or failed.
USER_END        Logs when a user ends session.

Note:

Mandatory Filters:

The HOST_EVENTS filter and its sub filter host-event are mandatory fields for auditd.
The minimum required CLI command to capture host events is shown below:

Ex: efa notification subscribers add-syslog-
relp --address 10.37.11.38:20514 --insecure --filter HOST_EVENTS --host-event auditd

Default Behavior:

If no key and record type are specified,
XCO will default to sending only USER_LOGIN type audit logs to the external server.

Ex: efa notification subscribers add-syslog-
relp --address 10.37.11.38:20514 --insecure --filter HOST_EVENTS --host-event auditd
```

- Enable a Syslog subscriber with HOST_EVENTS filter type, focusing on host events with Auditd event, using specific audit rule group key "logins" and record types USER_LOGIN and USER_END, and establishes a secure connection with an external server using a CA certificate.

```
(efa:extreme)extreme@tpvm:~$ (efa:user)user@pkumarpatra22041:~$ efa notification
subscribers add-syslog-relp --address "10.37.11.21:20514" --filter HOST_EVENTS --host-
event auditd --host-auditd-key logins --minimum-severity warning --host-auditd-record-
type USER_LOGIN,USER_END --cacert /path/to/ca-cert.pem
Successfully registered subscriber.

+-----------
+-------------------------------------------------------------------------------------+
| attribute |
value                                                                                |
+-----------
+-------------------------------------------------------------------------------------+
| id        |
```

```
17                                                                                |
+-----------
+-------------------------------------------------------------------------------------+
| handler   |
relp                                                                              |
+-----------
+-------------------------------------------------------------------------------------+
| endpoint
| 10.37.11.21:20514                                                                  |
+-----------
+-------------------------------------------------------------------------------------+
| config
| {"cacert":"<Certificate Content>","conn-timeout":10,"device-event                  |
|           |
":[],"filters":["HOST_EVENTS"],"host-auditd-key":["logins"],"host- auditd-record|
|           |
-type":["USER_LOGIN","USER_END"],"host-event":["auditd"],"insecure"                 |
|           | :false,"minimum-
severity":"warning","rfc5424":false}                                |
+-----------
+-------------------------------------------------------------------------------------+
```

- Enable a Syslog subscriber to capture various event types, including HOST_EVENTS, DEVICE_EVENTS, and APP_EVENTS. This configuration utilizes specific audit rule group keys "logins" and record types (USER_LOGIN and USER_END) for auditd events, as well as audit-security and audit-configuration for device events. The following CLI command enables logging for host and device events across multiple categories, transmitting the logs over an insecure channel:

```
(efa:extreme)extreme@tpvm:~$ efa notification subscribers add-syslog-relp --address
127.0.0.1:1601 --insecure --filter APP_ALERTS,DEVICE_EVENTS,HOST_EVENTS --device-event
audit-security,audit-configuration  --host-event auditd --host-auditd-key logins --
host-auditd-record-type USER_LOGIN,USER_END
Successfully registered subscriber.

+-----------
+-------------------------------------------------------------------------------------+
| attribute |
value                                                                             |
+-----------
+-------------------------------------------------------------------------------------+
| id        |
19                                                                                |
+-----------
+-------------------------------------------------------------------------------------+
| handler   |
relp                                                                              |
+-----------
+-------------------------------------------------------------------------------------+
| endpoint
| 127.0.0.1:1601                                                                     |
+-----------
+-------------------------------------------------------------------------------------+
| config
| {"cacert":"","conn-timeout":10,"device-event":["audit-security","audit-configura|
|           |
tion"],"filters":["APP_ALERTS","DEVICE_EVENTS","HOST_EVENTS"],"host-auditd-key":|
|           |
["logins"],"host-auditd-record-type":["USER_LOGIN","USER_END"],"host-event":    |
|           |
["auditd"],"insecure":true,"minimum-severity":"","rfc5424":false}                 |
+-----------
+-------------------------------------------------------------------------------------+
```

```
Notification Subscriber ID=19
-- Time Elapsed: 2.04493432s ---
```

- Enable a Webhook subscriber for various event types, including HOST_EVENTS, DEVICE_EVENTS, and APP_EVENTS. Specifically, it utilizes audit rule group keys "logins" and record types USER_LOGIN and USER_END for auditid events. For device events, it monitors audit-security and audit-configuration. The following CLI command ensures that logging is enabled for both host and device events across multiple categories, transmitted over an insecure channel:

```
(efa:extreme)extreme@tpvm:~$ efa notification subscribers add-https --url
https://127.0.0.1:5000 --username jarvis --password vision --insecure --filter
APP_ALERTS,DEVICE_EVENTS,HOST_EVENTS --device-event audit-security,audit-configuration
--host-event auditd --host-auditd-key logins --minimum-severity warning --host-auditd-
record-type USER_LOGIN,USER_END
Successfully registered subscriber.

+-----------
+-------------------------------------------------------------------------------------+
| attribute |
value                                                                                |
+-----------
+-------------------------------------------------------------------------------------+
| id        |
21                                                                                   |
+-----------
+-------------------------------------------------------------------------------------+
| handler   |
http                                                                                 |
+-----------
+-------------------------------------------------------------------------------------+
| endpoint
| https://127.0.0.1:5000                                                              |
+-----------
+-------------------------------------------------------------------------------------+
| config
| {"cacert":"","conn-timeout":10,"device-event":["audit-security","audit-configura|
|           |
tion"],"filters":["APP_ALERTS","DEVICE_EVENTS","HOST_EVENTS"],"host-auditd-key":|
|           |
["logins"],"host-auditd-record-type":["USER_LOGIN","USER_END"]                       |
|           |
"host-event":["auditd"],"insecure":true,"minimum-severity":"warning","password":|
|           |
"vision","username":"jarvis"}                                                        |
+-----------
+-------------------------------------------------------------------------------------+
Notification Subscriber ID=21
--- Time Elapsed: 2.042500359s ---
```

Notification Subscriber REST API Changes for Audit Logs Filtering

```
curl --location --request POST http://gonotification-service:8088/v1/notification/
subscribers \

--header 'Content-Type: application/json' \

--data-raw '{
            "username": "jarvis",
            "password": "vision",
            "insecure": true,
            "conn-timeout": 10,
```

```
            "filters": ["HOST_EVENTS"],
            "rfc5424":true,
            "host-event":["auditd"],
            "host-auditd-key ":["logins","time-change"],
            "host-auditd-record-type":[USER_LOGIN,LOGIN],
            "ca-cert":"",
            "minimum-severity":""
}'
## valid audit keys are
**"username", "password", "insecure", "cacert", "conn-timeout" "filters" "device-event"
"host-event" "minimum-severity" "host-auditd-key" and host-auditd-record-type **.

|-----------------------|-----------------------------------------------|------------|
|   Valid keys          |            Description                        | Handler    |
|-----------------------|-----------------------------------------------|------------|
|   username            | username of the webhook/http url              | http       |
|                       |                                               |            |
|   password            | password of the webhook/http url              | http       |
|                       |                                               |            |
|   insecure            | indicates if the connection is secure or not  | http & relp|
|                       |                                               |            |
|   ca-cert             | CA certificate for secure connection          | http & relp|
|                       |                                               |            |
|   filters             | list of filters based on which selected       | http & relp|
|                | notifications will be filtered to subscribers|           |
|                       |                                               |            |
|   device-event        | list of allowed device event notifications    | http & relp|
|                       |                                               |            |
|   minimum-severity    | minimum-severity-value                        | http & relp|
|                       |                                               |            |
|   conn-timeout        | connection timeout for http/relp in seconds   | http & relp|
|                       |                                               |            |
|   rfc5424             | enable RFC5424 message format for syslog      | relp       |
|                       |                                               |            |
|   host-event          | Indicate list of host event category          | http & relp|
|                       | (auditd)                                      |            |
|   host-auditd-key     | list of audit rule key group/category         | http & relp|
|                       |                                               |            |
|host-auditd-record-type| list of audit rule key group/category         | http & relp|
|-----------------------|-----------------------------------------------|------------|
```

Storage Limits and Log Rotations

Use the following auditd configuration for audit log storage limits and log rotations:

```
num_logs = 5
max_log_file = 100 // in megabytes
max_log_file_action = ROTATE
```

High Availability

Audit logs will be generated for each node, recording the IP address of the device (primary or standby) where the user logs in.

Logs will be forwarded from both active and standby nodes to an external server, providing a comprehensive audit log from both nodes. The auditd service on both nodes operates in an active-active configuration.

XCO and TPVM Upgrade

During the TPVM upgrade, the TPVM handles the upgrade of the auditd service. The hardening script includes default audit rules, which must be run manually. Additionally, XCO modifies auditd configurations to facilitate the forwarding of audit logs to an external server, thereby ensuring uninterrupted audit log forwarding.

The auditd rsyslog configuration will remain unchanged during the upgrade process; no modifications will be made.

If auditd rsyslog forwarding was enabled before the upgrade, it will remain enabled afterward. Likewise, if it was disabled prior to the upgrade, it will continue to be disabled post-upgrade.

Post XCO upgrade, the auditd and rsyslog services are restored to their previous state, ensuring they remain enabled and running.

Note
The host events are lost during the XCO and TPVM upgrade process.

Across all upgrade scenarios (XCO full upgrade, XCO incremental upgrade, TPVM full upgrade, and TPVM incremental upgrade), XCO will neither enforce nor remove the existing rsyslog configuration.

- If the rsyslog configuration is present before the XCO or TPVM upgrade, it will remain intact after the upgrade.
- If the rsyslog configuration is not present prior to the upgrade, it will not be introduced or added during the upgrade.

For XCO node replacement and single-node to multi-node expansion, the rsyslog configuration will be replicated from the active node to ensure consistency.

Note
For a seamless upgrade to XCO 4.0.2 or later, run the following Linux shell command to permanently delete the unwanted TLS configuration file (for example, 40-raslog-tls.conf) for rsyslog:"rm /etc/rsyslog.d/40-raslog-tls.conf"systemctl restart rsyslog:

```
"rm /etc/rsyslog.d/40-raslog-tls.conf"
systemctl restart rsyslog
```

Ensure that you run the Linux shell command on each node individually. In a multi-node deployment, this applies to both active and standby nodes.

Supportsave Contents and Logging Coverage

Supportsave includes RASLog and Notification service logs for real-time system health and alert visibility. XCO integrates a full set of auditd rules, configuration files (auditd.conf, rsyslog.conf) and log files for debugging purposes.

All operational logs from Monitor, System, and Installer services are now automatically bundled into the supportsave package and retained for troubleshooting and audit review.

Backup & Restore

Only the database will be backed up during the backup process, allowing for restoration of the same subscriber.

XCO Backup now captures auditd and rsyslog config files. During restore, you can apply these config files either on the same node or on a new node.

Validate Security Event Logs

About This Task
Follow this procedure to validate audit log on the external server.

Procedure

1. Setup external server.

   Set up and configure an external syslog or webhook server to receive logs. Users can choose a secure or insecure channel.
2. Register external log recipient server with XCO.

   Register the external server as an **efa notification subscriber** to receive audit logs. Choose any of the following options to add notification subscribers to XCO:

   a. Insecure syslog-relp
   b. Secure syslog-relp
   c. Insecure Webhook
   d. Secure Webhook

   The following example shows an insecure syslog relp CLI configuration:

```
efa:extreme)extreme@tpvm:~$ efa notification subscribers add-syslog-relp
--address 127.0.0.1:1601 --insecure --filter HOST_EVENTS
--host-event auditd
--host-auditd-key logins,time_change --minimum-severity warning
--host-auditd-record-type USER_LOGIN, USER_LOGOUT
Successfully registered subscriber.

+-----------+-----------------------------------------------------------------+
| attribute | value                                                           |
+-----------+-----------------------------------------------------------------+
| id        | 17                                                              |
+-----------+-----------------------------------------------------------------+
| handler   | relp                                                            |
+-----------+-----------------------------------------------------------------+
| endpoint  | 127.0.0.1:1601                                                  |
+-----------+-----------------------------------------------------------------+
| config    | {    "cacert":"",                                                 |
|           |    "conn-timeout":10                                            |
|           |    "filters":["HOST_EVENTS"],                                   |
|           |    "host-event":["auditd"],                                     |
|           |     "host-auditd-key ":[logins,time_change"],                   |
|           |     "host-auditd-record-type":[USER_LOGIN, USER_LOGOUT],
|           |     "insecure":true,                                            |
|           |    "minimum-severity":"warning"                                 |
|           |     "rfc5424":false                                             |
|           | }                                                               |
+-----------+-----------------------------------------------------------------+
```

```
Notification Subscriber ID=17
--- Time Elapsed: 2.171041915s ---
```

3. Trigger audit log events.

   Trigger login or logout audit events by logging in to TPVM.

4. Validate logs.

   Verify that the audit logs are received on the external server. If using RFC5424 for adding a notification subscriber, the log message will follow the standard format.

```
<PRI>VERSION TIMESTAMP HOSTNAME APP-NAME PROC-ID MSGID [STRUCTURED-DATA] AUDIT LOG MSG

node=127.0.1.1 type=USER_START msg=audit(1742293858.104:1365): pid=48734 uid=0
auid=1000 ses=16 subj=unconfined msg='op=PAM:session_open
grantors=pam_selinux,pam_loginuid,pam_keyinit,pam_permit,pam_umask,pam_unix,pam_systemd
,pam_mail,pam_limits,pam_env,pam_env,pam_selinux acct="user" exe="/usr/sbin/sshd"
hostname=134.141.202.209 addr=134.141.202.209 terminal=ssh res=success' UID="root"
AUID="user"
```

   For example,

```
<14>1 2025-03-18T10:30:58.110734Z 10.32.85.28 audit 48734 AUDITD_EVENTS -
node=127.0.1.1 type=USER_START msg=audit(1742293858.104:1365): pid=48734 uid=0
auid=1000 ses=16 subj=unconfined msg='op=PAM:session_open
grantors=pam_selinux,pam_loginuid,pam_keyinit,pam_permit,pam_umask,pam_unix,pam_systemd
,pam_mail,pam_limits,pam_env,pam_env,pam_selinux acct="user" exe="/usr/sbin/sshd"
hostname=134.141.202.209 addr=134.141.202.209 terminal=ssh res=success' UID="root"
AUID="user".
```

Enable or Disable Host OS Auditd Log Forwarding to XCO

You can enable or disable the forwarding of Host OS auditd event logs to ExtremeCloud Orchestrator (XCO) and the associated user notifications at any point during deployment using either CLI commands or REST API calls.

About This Task

This capability allows you to temporarily disable the Host OS auditd event logs forwarding to XCO during critical maintenance windows, sensitive operations or troubleshooting without losing local logs, then re-enable seamlessly once system stability is restored. No impact to switch rsyslog. All logs persist for audit and troubleshooting.

Note

- Local logging continues even when the XCO auditd feature is disabled
- Access the logs remain on the node at
  - /var/log/audit/audit.log
  - /var/log/syslog:
- All auditd-related logs (forwarder, processor, indexer) are also available in
  - Installer logs
  - Monitor service logs
  - Inventory service logs
- SLX or OS ONE switch rsyslog events are unaffected by this toggle. They operate independently.

Follow the procedure to enable or disable the Auditd feature.

For detailed command syntax, parameter descriptions, and REST API endpoints, refer to the *ExtremeCloud Orchestrator Command Reference Guide*.

Procedure

1. To enable the auditd rsyslog auditd config on the node, run the following command:

```
efa system feature update --host-auditd-rsyslog Enable
 Feature Setting Updated Successful
--- Time Elapsed: 19.98308142s ---
```

2. To disable the auditd rsyslog auditd config on the node, run the following command:

```
efa system feature update --host-auditd-rsyslog Disable
Feature Setting Updated Successfully
--- Time Elapsed: 19.98308142s ---
```

3. To view all feature flags and status, run the following command:

```
efa system feature show
+-------------------------------------------+-------------+
|                FEATURE NAME               |   STATUS    |
+-------------------------------------------+-------------+
| Inflight Transaction Auto Recovery        | Enabled     |
+-------------------------------------------+-------------+
| Tenant Api Concurrency                    | Enabled     |
+-------------------------------------------+-------------+
| Tenant EPG Api Bulking                    | Disabled    |
+-------------------------------------------+-------------+
| Host AuditD RSyslog                       | Enabled     |
+-------------------------------------------+-------------+
--- Time Elapsed: 110.860783ms ---
```

Rsyslog Configuration for Default Auditd Logging Filter

XCO adds filter rules into the host rsyslog service to restrict which auditd events are forwarded to the default logging flow (for example, to XCO or remote syslog servers)

XCO can dynamically configure rsyslog to selectively forward security-relevant auditd events via TLS, improving log hygiene and system performance without affecting local logging. This means:

- Only filtered audit events are sent to the configured remote syslog server.
- All other auditd messages remain local (in /var/log/audit/audit.log, /var/log/syslog) and are not forwarded.
- Reduced bandwidth and storage overhead while preserving critical security logs.
- Reduced noise, improve performance, and focus on critical audit trails

The following is an example of XCO-generated rsyslog config:

```
$$$>> cat /etc/rsyslog.d/40-raslog-tls.conf

# Enable encrypted communication
module(load="lmnsd_gtls")

# Configure TLS for TCP
$DefaultNetstreamDriver gtls
$ActionSendStreamDriverMode 1
$ActionSendStreamDriverAuthMode anon
$DefaultNetstreamDriverCAFile /opt/efadata/certs/ca/extreme-ca-cert.pem

template(name="SyslogProtocolRFC5424_Format" type="string"
```

```
     string="<%PRI%>1 %timegenerated:::date-rfc3339% 10.32.85.28 audit %procid%
AUDITD_EVENTS - %msg%\n"
)

template(name="SyslogProtocolRFC5424_Format" type="string"
     string="<%PRI%>1 %timegenerated:::date-rfc3339% 10.32.85.28 audit %procid%
AUDITD_EVENTS - %msg%\n"
)

if $programname == 'audispd' or $programname == 'audisp-syslog' then {
    if (
        $msg contains "USER_LOGIN" or
        $msg contains "USER_LOGOUT" or
        $msg contains "USER_AUTH" or
        $msg contains "USER_ACCT" or
        $msg contains "USER_START" or
        $msg contains "LOGIN" or
        $msg contains "USER_END" or
        $msg contains "CRED_ACQ" or
        $msg contains "CRED_DISP" or
        $msg contains "CRED_REFR" or
        $msg contains "logins"
     ) then {
*.* @@10.32.85.37:6514;SyslogProtocolRFC5424_Format
        stop
    }
}
```

Configure Additional Auditd Record Types and Keys

You can extend the default auditd filter set by defining custom auditd filters (custom record types or keys) in the rsyslog configuration file, which enables control over which auditd events are forwarded to XCO.

About This Task

Follow this procedure to receive auditd events for "file or directory paths accessed during syscalls" (record type: PATH).

Note

Ensure to explicitly list all auditd record types or keys configured in XCO notification subscribers in the rsyslog configuration of the node. Failure to do so (if a record type is missing from rsyslog) results in those events not being forwarded to XCO, even if they're generated locally.

Procedure

1. Add Notification Subscriber with desired `record type` and `auditd key`.

   Run the following command for XCO to create a subscriber that filters for auditd events with auditd record type of "PATH":

```
efa notification subscribers add-syslog-relp \
    --address 10.37.12.20514 \
    --insecure \
    --filter HOST_EVENTS \
    --host-event auditd \
    --host-auditd-record-type PATH
```

2. Update the `rsyslog` config file to include PATH (same record type in filter).

   Edit the `/etc/rsyslog.d/40-raslog-tls.conf` (or equivalent) file and add PATH to the message filter to ensure that auditd events are forwarded correctly to the XCO:

```
$$$>> cat /etc/rsyslog.d/40-raslog-tls.conf

# Enable encrypted communication
module(load="lmnsd_gtls")

# Configure TLS for TCP
$DefaultNetstreamDriver gtls
$ActionSendStreamDriverMode 1
$ActionSendStreamDriverAuthMode anon
$DefaultNetstreamDriverCAFile /opt/efadata/certs/ca/extreme-ca-cert.pem

template(name="SyslogProtocolRFC5424_Format" type="string"
     string="<%PRI%>1 %timegenerated:::date-rfc3339% 10.32.85.28 audit %procid%
AUDITD_EVENTS - %msg%\n"
)


if $programname == 'audispd' or $programname == 'audisp-syslog' then {
    if (
        $msg contains "USER_LOGIN" or
        $msg contains "USER_LOGOUT" or
        $msg contains "USER_AUTH" or
        $msg contains "USER_ACCT" or
        $msg contains "USER_START" or
        $msg contains "LOGIN" or
        $msg contains "USER_END" or
        $msg contains "CRED_ACQ" or
        $msg contains "CRED_DISP" or
        $msg contains "CRED_REFR" or
        $msg contains "logins" or
        $msg contains "PATH"
     ) then {
*.* @@10.32.85.37:6514;SyslogProtocolRFC5424_Format
        stop
    }
}
```

   Note
   To ensures rsyslog correctly matches and forwards them to the remote destination, always add new record types/keys to the `if` condition block.

3. Verify auditd rule exists for PATH.

   Ensure that an auditd rule exists to generate events for the PATH record type (or any custom key). If missing, create it.

   Note
   The default hardening script may not cover all record types. You must manually add rules for custom types or keys.

4. Restart rsyslog service.

   Apply the changes by restarting rsyslog:

```
systemctl restart rsyslog.service
systemctl status rsyslog.service
```

5. Verify status is active (running) and no errors appear.

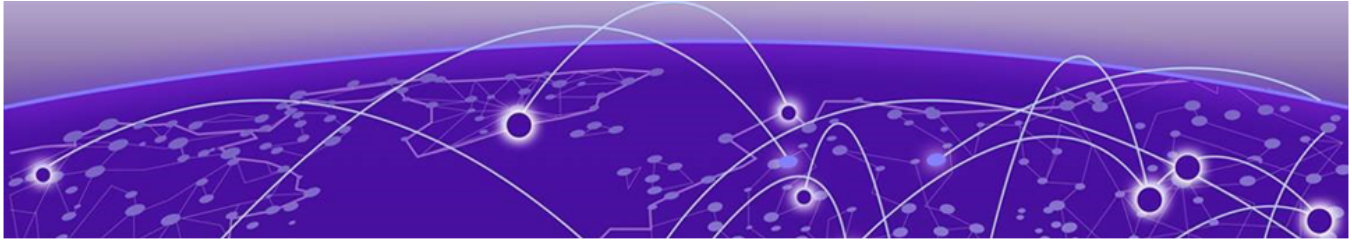

Unified Health and Fault Management

Learn about managing alerts, alarms, and health status of managed entities in XCO.

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 and corresponding alarm raised by the Fault Management system.

Fabric Health provides a centralized view of the operational and configuration state of the fabric. It is computed and maintained by the Fabric Service and reflects whether the fabric is correctly configured, structurally complete, and operationally functional.

Fabric Health is not manually configured. Instead, it is continuously derived from authoritative system inputs and updated as changes occur.

Health states are evaluated using severity precedence:

Critical > Degraded > Healthy

The most severe condition determines the final reported health state.

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 and corresponding alarm 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.

Fabric Health Composition

Fabric Health is derived from both device-level and fabric-level contributors.

Device Health

Configuration State Health

Represents correctness of intent realization:

- Device provisioning status
- Configuration generation and application
- Detection of configuration drift

Operational State Health

Represents runtime behavior:

- Physical topology (interfaces, LLDP, link status)
- Underlay health (routing protocols, BGP sessions)
- Overlay health (VXLAN tunnels, evaluated only when tenant intent exists)
- Cluster and multihoming status (MCT/MLAG)

A device is considered healthy only when both configuration and operational states are healthy.

Fabric-Level Health

Fabric Health also includes the following:

- Physical Topology Health – Validates fabric structure, device roles, and topology integrity
- Lifecycle/Status Health – Tracks outcomes of operations such as configuration, upgrade, and migration

Final Health Calculation

The overall Fabric Health is computed using:

- Worst-case device health
- Fabric physical topology health
- Fabric lifecycle/status health

Data Sources and Authority Model

Fabric Health relies on clearly defined ownership across services.

Inventory Service

- Acts as the authoritative source for device configuration and runtime state
- Collects and normalizes telemetry
- Owns physical, underlay, overlay runtime, and cluster state

Note
The Fabric Service does not directly interact with network devices.

Tenant Service

- Defines overlay intent (EPGs, VNIs, endpoint relationships)
- Publishes overlay intent events when tenant configurations change
- Ensures overlay health is evaluated only when relevant intent exists

Fabric Service

- Computes, persists, and exposes Fabric Health
- Consumes inputs from Inventory and Tenant services
- Publishes health via APIs, streaming events (SSE), and alerts

Health Data Collection and Reconciliation

Fabric Health uses a hybrid model to ensure timely updates and eventual consistency.

Event-Driven Updates (Primary Mechanism)

- Devices generate RAS logs for state changes
- Inventory processes logs and performs scoped data collection
- Events are published to the Fabric Service
- Fabric Health is recomputed based on updates

A debounce mechanism aggregates events to prevent excessive updates. Under sustained activity, updates may be delayed up to 15 minutes.

Configuration Tracking Number (CTN)

CTN is used to control configuration discovery:

- CTN match → Configuration unchanged → Minimal discovery
- CTN mismatch → Potential drift → Deep discovery

This ensures controlled and efficient data collection.

Periodic and On-Demand Reconciliation

To ensure consistency, periodic and manual mechanisms are used:

- Cron-based reconciliation – Every 30 minutes
- Inventory timer-based collection – Every hour
- Service restart reconciliation – Pulls latest state after restart
- Manual trigger – Forces recomputation for validation

These mechanisms ensure convergence even if events are delayed or missed

Health Reconciliation Mechanisms

Mechanism	Trigger	Purpose
Event-driven	RAS logs	Near real-time updates
Scheduled reconciliation	Every 30 minutes	Recover missed updates
Inventory timer	Every hour	Telemetry consistency
Service restart	Restart event	Restore system state
Manual trigger	Operator action	On-demand validation

Known Behavioral Considerations and Enhancements

Known Behavioral Considerations

Event Sequencing Noise

Certain link event sequences (for example, link add/remove ordering) may appear as configuration drift during normal operations such as:

- Port flaps
- Device reloads
- Administrative interface changes

These conditions can temporarily impact reported health despite no actual configuration issues.

Transient State Conditions

Temporary inconsistencies (for example, multiple IP assignments during convergence) may be interpreted as drift and impact health.

Aggregation Trade-offs

- Event aggregation may delay health updates during high activity
- In-memory aggregation context may be lost during failover, requiring reconciliation

Observability Gaps

If upstream services are unreachable:

- Health state may not reflect current conditions
- Earlier behavior marked these as unhealthy; updated behavior preserves last known state

Enhancements

Health Stability Enhancements

- Event sequencing issues are handled as operational changes
- Configuration state is preserved unless actual drift is detected
- Reduction in false health degradation events

Improved Handling of Unobservable States

- Health is not degraded when dependent services are unreachable
- Last known state is retained to avoid false alarms

Service Reliability Improvements

- Health endpoints validate dependencies
- Unhealthy services are automatically restarted
- Enhances system resiliency and recovery

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, 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.

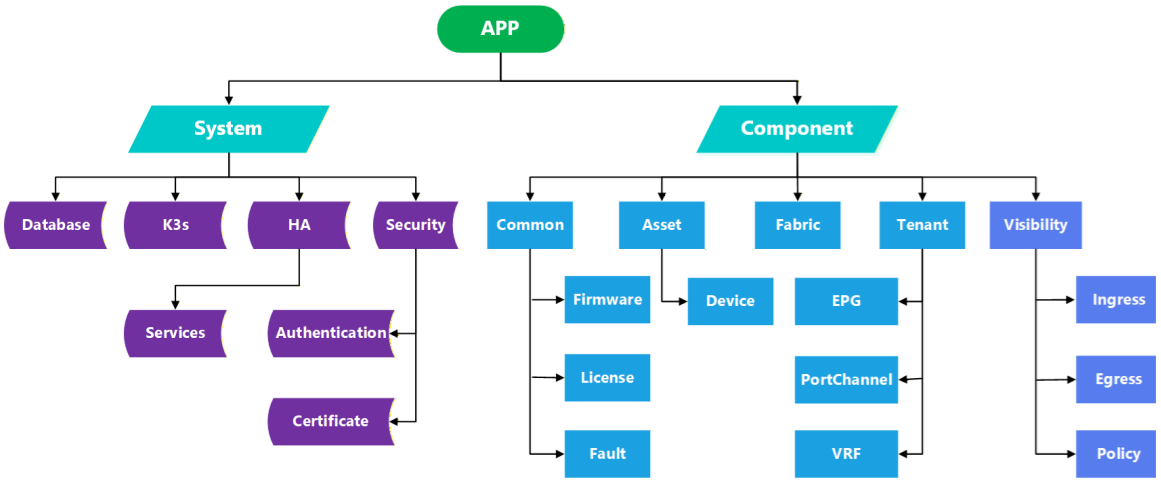

Figure 56: XCO Health Resource Hierarchy

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	3.8.0	Software Version (of XCO host)	
Resource	alert@1916	/App/System/Security/Certificate?type=app_server_certificate	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=”3.8.0”]
   [alert@1916
    resource=”/App/System/Security/Certificate?type=app_server_certificate”
    alertId=”31000”
    cause=”keyExpired”
    type=”securityServiceOrMechanismViolation”
    severity=”warning”]
   [alertData@1916
    type="app_server_certificate"
    expiry_data="Sep 12 10:00:45 2023 GMT"]
   BOMThe application server certificate on the application will expire soon on “Sep 12
10:00:45 2023 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
SystemBackupRestoreFailureAlert	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 from the sequence 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_certificate
					default_intermediate_ca
					default_root_ca
					third_party_ca
					k3s_server_certificate
					k3s_ca
					jwt_certificate
DeviceCertificateExpiryNoticeAlert	31001	Warning	/App/System/Security/Certificate	device_ip	IP of the affected device
				type	https_server_certification
					syslog_ca
					jwt_verifier
CertificateExpiredAlert	31002	Critical	/App/System/Security/Certificate	type	app_server_certificate
					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_certificate
					default_intermediate_ca
					default_root_ca
					third_party_ca
					jwt_certificate
DeviceCertificateRenewalAlert	31005	Info	/App/System/Security/Certificate	device_ip	IP of the device
				type	https_server_certification
					jwt_verifier
CertificateUnreadableAlert	31006	Warning	/App/System/Security/Certificate	type	app_server_certificate
					default_intermediate_ca
					default_root_ca
					third_party_ca
					k3s_server_certificate
					k3s_ca
					jwt_certificate
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_certification
					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

Device Link Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
PortFlapAlert	31600	Warning	/App/Component/Asset/Device/PortFlap	device_ip	
				port	
PortFlapClearAlert	31601	Info	/App/Component/Asset/Device/PortFlap	device_ip	
				port	

Fabric Health Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
FabricCreatedAlert	31700	Info	/App/Component/Fabric	fabric_name	
FabricDeletedAlert	31701	Info	/App/Component/Fabric	fabric_name	
FabricDeviceAddedAlert	31702	Info	/App/Component/Fabric/Device	fabric_name	
FabricDeviceRemovedAlert	31703	Info	/App/Component/Fabric/Device	fabric_name	
FabricStateDegradedAlert	31704	Major, Critical	/App/Component/Fabric/State	fabric_name	
FabricStateHealthyAlert	31705	Info	/App/Component/Fabric/State	fabric_name	
FabricPhysicalTopologyDegradedAlert	31706	Major, Critical	/App/Component/Fabric/Topology/Physical	fabric_name	
FabricPhysicalTopologyHealthyAlert	31707	Info	/App/Component/Fabric/Topology/Physical	fabric_name	

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
FabricDeviceAppStateDegradedAlert	31708	Major, Critical	App/Component/Fabric/Device/Configuration/AppState	fabric_name	
				device_ip	
FabricDeviceAppStateHealthyAlert	31709	Info	App/Component/Fabric/Device/Configuration/AppState	fabric_name	
				device_ip	
FabricDeviceProvisioningStateDegradedAlert	31710	Major, Critical	App/Component/Fabric/Device/Configuration/DevState	fabric_name	
				device_ip	
FabricDeviceProvisioningStateHealthyAlert	31711	Info	App/Component/Fabric/Device/Configuration/DevState	fabric_name	
				device_ip	
FabricDeviceMctDegradedAlert	31712	Major, Critical	/App/Component/Fabric/Device/Operational/Mct	fabric_name	
				device_ip	
FabricDeviceMctHealthyAlert	31713	Info	/App/Component/Fabric/Device/Operational/Mct	fabric_name	
				device_ip	
FabricDevicePhysicalTopologyDegradedAlert	31714	Major, Critical	/App/Component/Fabric/Device/Operational/Topology/Physical	fabric_name	
				device_ip	
FabricDevicePhysicalTopologyHealthyAlert	31715	Info	/App/Component/Fabric/Device/Operational/Topology/Physical	fabric_name	
				device_ip	
FabricDeviceUnderlayTopologyDegradedAlert	31716	Major, Critical	/App/Component/Fabric/Device/Operational/Topology/Underlay	fabric_name	
				device_ip	
FabricDeviceUnderlayTopologyHealthyAlert	31717	Info	/App/Component/Fabric/Device/Operational/Topology/Underlay	fabric_name	
				device_ip	
FabricDeviceOverlayTopologyDegradedAlert	31718	Major, Critical	/App/Component/Fabric/Device/Operational/Topology/Overlay	fabric_name	
				device_ip	
FabricDeviceOverlayTopologyHealthyAlert	31719	Info	/App/Component/Fabric/Device/Operational/Topology/Overlay	fabric_name	
				device_ip	

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
FabricHealthDegradedAlert	31799	Major, Critical	/App/Component/Fabric	fabric_name	
FabricHealthRestoredAlert	31800	Info	/App/Component/Fabric	fabric_name	

High Availability Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
HAServiceNonRedundantAlert	31050	Minor	/App/System/HA/Nodes/Node		
HAServiceFullyRedundantAlert	31051	Info	/App/System/HA/Nodes/Node		
HAServiceNewActiveAlert	31052	Major	/App/System/HA/Nodes/Node		
ServiceDegradedAlert	31053	Warning	/App/System/HA/Nodes/Services		
ServiceRestoredAlert	31054	Info	App/System/HA/Nodes/Services		

License Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
LicenseExpiryThresholdAlert	31400	Warning Minor Major	*/App/System/Security/License*	AID	
LicenseExpiredAlert	31401	Critical	*/App/System/Security/License*	AID	
LicenseExpiryClearAlert	31402	Info	*/App/System/Security/License*	AID	

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

Password Expiry Alerts Inventory

Alert Name	Alert ID	Severity	Resource	Query Params	Param Values
DevicePasswordExpiryThresholdAlert	35100	Warning	/App/Component/Asset/Device/Password	device_ip	IP address of SLX device
				username	User name
DevicePasswordExpiredAlert	35101	Minor	/App/Component/Asset/Device/Password	device_ip	IP address of SLX device
				username	User name
DevicePasswordExpiryClearAlert	35102	Info	/App/Component/Asset/Device/Password	device_ip	IP address of SLX device
				username	User name

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.

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.x.x.x” enterpriseId=”1916” software=”XCO” swVersion=”3.4.0”] [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=”3.4.0”] [alert@1916 resource=”” alertId=”31070” cause=”operationInitiated” type=”communicationsAlarm” severity=”info”] [alertData@1916 user=”extreme” version=”3.4.0-backup” backup_type=”FabricName” fabric_name=”nonclos” fabric_all=”” device_ips=””] BOMSystem backup type "FabricName" initiated for fabric "nonclos" for efa version "3.4.0-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=”3.4.0”] [alert@1916 resource=”” alertId=”31071” cause=”operationSuccess” type=”communicationsAlarm” severity=”info”] [alertData@1916 user=”extreme” version=”3.4.0-backup” backup_type=”FabricAll” fabric_name=”” fabric_all=true device_ips=”” filename="/apps/efa_logs/backup/XCO-3.4.0-backup-2022-12-07T19-05-31.719.tar”] BOMSystem backup type "FabricAll" completed successfully for efa version "3.4.0-backup" by user "extreme" located at "/apps/efa_logs/backup/XCO-3.4.0-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=”3.4.0”] [alert@1916 resource=”” alertId=”31072” cause=”operationFailure” type=”processingErrorAlarm” severity=”major”] [alertData@1916 user=”extreme” version=”3.4.0-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 "3.4.0-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="3.4.0"] [alert@1916 resource="" alertId="31075" cause="operationInitiated" type="communicationsAlarm" severity="info"] [alertData@1916 user="extreme" version="3.4.0-backup" filename=" XCO-3.4.0-backup-2022-12-05T15-48-54.920.tar"] BOMSystem restore initiated with "XCO-3.4.0-backup-2022-12-05T15-48-54.920.tar" on XCO version "3.4.0-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=”3.4.0”]`
	`[alert@1916`
	`resource=””`
	`alertId=”31076”`
	`cause=”operationSuccess”`
	`type=”communicationsAlarm”`
	`severity=”info”]`
	`[alertData@1916`
	`user=”extreme”`
	`version=”3.4.0-backup”`
	`filename=” XCO-3.4.0-backup-2022-12-05T15-48-54.920.tar”]`
	`BOMSystem restore completed successfully with "XCO-3.4.0-backup-2022-12-05T15-48-54.920.tar" on XCO version "3.4.0-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="3.4.0"]`
	`[alert@1916`
	`resource=""`
	`alertId="31077"`
	`cause="operationFailure"`
	`type="processingErrorAlarm"`
	`severity="major"]`
	`[alertData@1916`
	`user=extreme""`
	`version="3.4.0-backup"`
	`filename=" XCO-3.4.0-backup-2022-12-05T15-48-54.920.tar"`
	`error="reason for failure"]`
	`BOMSystem restore failed with "XCO-3.4.0-backup-2022-12-05T15-48-54.920.tar" on XCO version "3.4.0-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="3.4.0"]`
	`[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_certificate ◦ Default Intermediate CA: default_intermediate_ca ◦ Default Root CA: default_root_ca ◦ Third-Party CA: third_party_ca ◦ K3s Server Certificate: k3s_server_certificate ◦ K3s CA: k3s_ca ◦ JWT Certificate: jwt_certificate 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="3.4.0"] [alert@1916 resource="/App/System/Security/Certificate?type=app_server_certificate" alertId="31000" cause="keyExpired" type="securityServiceOrMechanismViolation" severity="warning"] [alertData@1916 type="app_server_certificate" 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_certificate HQI { Color: Yellow Value: 2 } 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_certification ◦ 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="3.4.0"] [alert@1916 resource="/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_certification" alertId="31001" cause="keyExpired" type="securityServiceOrMechanismViolation" severity="warning"] [alertData@1916 device_ip="10.10.10.1" type="https_server_certification" 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_certification HQI { Color: Yellow Value: 2 } 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_certificate ◦ 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="3.4.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?type=app_server_certificate"`
	`alertId="31002"`
	`cause="keyExpired"`
	`type="securityServiceOrMechanismViolation"`
	`severity="critical"]`
	`[alertData@1916`
	`type="app_server_certificate"`
	`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_certificate`
	`HQI {`
	`Color: Black`
	`Value: 5`
	`}`
	`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=”3.4.0”]`
	`[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: 5`
	`}`
	`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_certificate ◦ Default Intermediate CA: default_intermediate_ca ◦ Default Root CA: default_root_ca ◦ Third-Party CA: third_party_ca ◦ JWT Certificate: jwt_certificate ◦ K3s Server Certificate: k3s_server_certificate ◦ 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=”3.4.0”]`
	`[alert@1916`
	`resource=”/App/System/Security/Certificate?type=app_server_certificate”`
	`alertId=”31004”`
	`cause=”keyGenerated”`
	`type=”securityServiceOrMechanismViolation”`
	`severity=”warning”]`
	`[alertData@1916`
	`type=”app_server_certificate”]`
	`BOMThe App Server Certificate on the application has bee renewed.`

Health Response	Response
	`{`
	`Resource: /App/System/Security/Certificate?type=app_server_certificate`
	`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_certification • 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=”3.4.0”]`
	`[alert@1916`
	`resource=”/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_certification”`
	`alertId=”31005”`
	`cause=”keyGenerated”`
	`type=”securityServiceOrMechanismViolation”`
	`severity=”info”]`
	`[alertData@1916`
	`device_iP=”10.10.10.1”`
	`type=”https_server_certification”]`
	`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_certification`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The HTTPS Server Certificate on the device 10.10.10.1 has been renewed.`
	`}`

XCO Certificate Unreadable Alert

31006	XCO Certificate Unreadable Alert
Description	Send an alert when XCO is unable to read the certificate.
Preconditions	Certificate Manager Component (Monitor & Auth Service) has system default settings that are NOT user-configurable. • Polling frequency for certificate expiry notice: **daily** • Monitors the following XCO Certificate Types: ◦ App Server Certificate (of XCO): app_server_certificate ◦ Default Intermediate CA: default_intermediate_ca ◦ Default Root CA: default_root_ca ◦ Third-Party CA: third_party_ca ◦ K3s Server Certificate: k3s_server_certificate ◦ K3s CA: k3s_ca ◦ JWT Certificate: jwt_certificate ◦ GlusterFS certficate: glusterfs_certficate ◦ Galera Certificate: galera_certificate The "DeviceCertificateUnreadableAlert" event notification is sent out daily with error message when XCO is unable to read a certificate of a particular type. The fault engine will process this event.

31006	XCO Certificate Unreadable Alert

Requirements	Alert shows the following data:
	• Certificate Type
	• Error
	
	The following example shows an alert when XCO is unable to read a certificate:
	`Syslog RFC-5424 Example:`
	`<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="3.8.0"]□□`
	`□□ [alert@1916`
	`resource="/App/System/Security/Certificate?type=app_server_certificate"`
	`alertId="31006"`
	`cause="keyExpired"`
	`type="securityServiceOrMechanismViolation"□`
	`severity="warning"]□`
	`□□ [alertData@1916□□`
	`□□ type="app_server_certificate"□□`
	`□□ error="Unable to read the expiration date of certificate"]□`
	`□□ BOMUnable to read app_server_certificate on the application due to Unable to read the expiration date of certificate.`

Health Response	Response
	`{`
	`Resource: /App/System/Security/Certificate?`
	`type=app_server_certificate`
	`HQI {`
	`Color: Yellow`
	`Value: 2`
	`}`
	`StatusText: Unable to read app_server_certificate on the application due to Unable to read the expiration date of certificate".□`
	`}`

XCO Device Certificate Unreadable Alert

31007	XCO Device Certificate Unreadable Alert
Description	Send an alert when XCO is unable to read the device certificate.
Preconditions	Certificate Manager Component (Monitor & Auth Service) has system default settings that are NOT user-configurable. • Polling frequency for certificate expiry notice: **daily** • Monitors the following XCO Certificate Types: ◦ HTTPS Server certificate: https_server_certficate ◦ JWT Verifier: jwt_verifier ◦ Syslog CA: syslog_ca The "DeviceCertificateUnreadableAlert" event notification is sent out daily with error message when XCO is unable to read a certificate of a particular type on a particular device. The fault engine will process this event.

Requirements	Alert shows the following data:
	• Certificate Type
	• Device IP
	• Error
	The following example shows an alert when when XCO is unable to read the device certificate:
	`Syslog RFC-5424 Example:`
	`<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="3.8.0"]`
	`[alert@1916`
	`resource="/App/System/Security/Certificate?type=https_server_certificate&device_ip=10.20.30.40"`
	`alertId="31007"`
	`cause="keyExpired"`
	`type="securityServiceOrMechanismViolation"`
	`severity="warning"]`
	`[alertData@1916`
	`type="app_server_certificate"`
	`device_ip="10.20.30.40"`
	`error="Unable to read the certificate"]`
	`BOMUnable to read https_server_certificate on the device 10.20.40.40 due to Certificate is not available.`

Health Response	Response
	`{`
	`Resource: /App/System/Security/Certificate?type=https_server_certificate&device_ip=10.20.30.40`
	`HQI {`
	`Color: Yellow`
	`Value: 2`
	`}`
	`StatusText: Unable to read https_server_certificate on the device 10.20.30.40 due to Certificate is not available".`
	`}`

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_certification ◦ 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="3.4.0"] [alert@1916 resource="/App/System/Security/Certificate?device_ip=10.10.10.1&type=https_server_certification" alertId="31008" cause="configRemoved" type="securityServiceOrMechanismViolation" severity="info"] [alertData@1916 device_ip="10.10.10.1" type="https_server_certification"] 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_certification 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=”3.4.0”] [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: 4 } 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="3.4.0"] [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="3.4.0"] [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.
	}

Device Link Alerts

Use the information in the following tables to learn about all possible fabric health alerts in detail that are raised by Fault Management.

Table 25: Port Flap Alert

31600	Port Flap Alert
Description	Send an alert when there is continuous port flap on a registered fabric device port.
Preconditions	The SFP is faulty or there are other hardware issues. The number of admin ups is greater than configured threshold within a specified time interval. The threshold value for number of admin ups are greater than 5 and the time interval is 30 seconds

Table 25: Port Flap Alert (continued)

31600	Port Flap Alert

Requirements	Alert shows the following data:
	• Device IP
	• Port Name
	The following is an example of a port flap alert:
	`<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.8.0"]`
	`[alert@1916`
	`resource=""`
	`alertId="31600"`
	`cause="excessiveErrorRate"`
	`type="communicationsAlarm"`
	`severity="warning"]`
	`[alertData@1916`
	`device_ip="10.x.x.x"`
	`port="Ethernet 0/1"]`
	`BOM The Port flap detected for device 10.x.x.x, port ethernet 0/1`

Health Response	*Response*
	`{`
	`Resource: /App/Component/Asset/Device?PortFlap?device_ip=1.1.1.1&port=Ethernet0/1`
	`HQI {`
	`Color: Yellow`
	`Value: 2`
	`}`
	`Status Text: Port Ethernet0/1 on device 10.x.x.x is flapping continuously.`
	`}`

Table 26: PortFlap Clear Alert

31601	PortFlap Clear Alert
Description	Send an alert when port flap stops on any registered device port.
Preconditions	After XCO raises Port flap alert, it will send Port flap clear alert if the flapping stops on that device port.

Table 26: PortFlap Clear Alert (continued)

31601	PortFlap Clear Alert

Requirements	Alert shows the following data:
	• Device IP
	• Port Name
	The following example shows a Port flap clear alert when the flapping stops on that device port:
	`<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=”3.8.0”]`
	`[alert@1916`
	`resource=”/App/Component/Asset/Device?device_ip=1.1.1.1&port=Ethernet0/1”`
	`alertId=”31601”`
	`cause=” operationSuccess”`
	`type=” communicationsAlarm”`
	`severity=”info”]`
	`[alertData@1916`
	`deviceip=”1.1.1.1”`
	`port=”Ethernet 0/1”]`
	`BOMThe Port flapping cleared for device1.1.1.1 on port ethernet 0/1`

Health Response	Response
	`{`
	`Resource: /App/Component/Asset/Device?`
	`device_ip=1.1.1.1&port=Ethernet0/1`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: Port flapping cleared for device1.1.1.1 on port ethernet 0/1`
	`}`

Fabric Health Alerts

Use the information in the following tables to learn about all possible fabric health alerts in detail that are raised by Fault Management.

Table 27: Managed Fabric Physical Topology Degraded Notice

31706	Managed Fabric Physical Topology Degraded Notice
Description	Send an alert when the fabric level physical topology health is changed from to Green to Red.
Preconditions	Fabric is created and devices are added in XCO. The severity for physical topology errors is **Major**. The fabric services generate the alerts when the following conditions are not met: • Fabric level physical topology validations for non-Clos fabric: 1. Each rack must contain two devices. • Fabric level physical topology validations for Clos fabric: 1. Stage 3 fabric must contain at least one leaf or border leaf device and spine device. 2. Stage 5 fabric must contain at least one leaf or border leaf device and super-spine devices

Table 27: Managed Fabric Physical Topology Degraded Notice (continued)

31706	Managed Fabric Physical Topology Degraded Notice
Requirements	Alert shows the following data: • Fabric Name • Fabric Health Info The following example shows an alert when a fabric level physical topology health is degraded: `<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric/Topology/Physical?fabric_name=fb" alertId="31706" cause="missingSpines" type="fabricService" severity="major"] [alertData@1916 fabric_name="fb" fabric_health_info="{ Missing_spines: true} }"] BOM The fabric "fb" is missing spines.`
Health Response	Response `{ Resource: /App/Component/Fabric/Topology/Physical?fabric_name =fb HQI { Color: Red Value: 4 } StatusText: The fabric "fb" is missing spines }`

Table 28: Managed Fabric Physical Topology Healthy Notice

31707	Managed Fabric Physical Topology Healthy Notice
Description	Send an alert when the fabric level physical topology health is changed from Red to Green.
Preconditions	Fabric is created and devices are added in XCO. The fabric services generate the alerts when the following conditions are met: • Fabric level physical topology validations for non-Clos fabric: 1. Each rack must contain two devices. • Fabric level physical topology validations for Clos fabric: 1. Stage 3 fabric must contain at least one leaf or border leaf device and spine device. 2. Stage 5 fabric must contain at least one leaf or border leaf device and super-spine devices

Table 28: Managed Fabric Physical Topology Healthy Notice (continued)

31707	Managed Fabric Physical Topology Healthy Notice

Requirements	Alert shows the following data:
	• Fabric Name
	• Fabric Health Info
	
	The following example shows an alert when a fabric level physical topology is in healthy condition:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916 resource="/App/Component/Fabric/Topology/Physical?fabric_name=fb"`
	`alertId="31707"`
	`cause="fabricPhysicalTopologyHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`fabric_health_info=""`
	`"`
	`]`
	`BOM The fabric "fb" has fabric physical topology in healthy state`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Topology/Physical?fabric_name =fb`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The fabric "fb" has fabric physical topology in healthy state.`
	`}`

Table 29: Managed Fabric Device Appstate Degraded Notice

31708	Managed Fabric Device Appstate Degraded Notice
Description	Send an alert when fabric device app state health is changed from Green to Red or Black.
Preconditions	Fabric is created and devices are added in XCO. Application state changes based on adding devices, configuring success or failure, drift in configurations. The alert has the following app states and the corresponding severity:

Table 29: Managed Fabric Device Appstate Degraded Notice

State	Severity	Description
cfg-ready	Major	Device configurations are ready to be pushed to device.
cfg-refreshed	Major	There is drift in configurations between switch and XCO intended configurations.
cfg-error	Critical	Configuration errors come before pushing configurations to the switch (for example, adding three leaf devices with links between all of them).
cfg-refresh-error	Major	Configuration failure like device reload state or missing links between device.
device-remove-failed	Critical	Remove device from fabric failed

Table 29: Managed Fabric Device Appstate Degraded Notice (continued)

31708	Managed Fabric Device Appstate Degraded Notice

Requirements	Alert shows the following data:
	• Fabric Name
	• Device IP
	• Fabric Health Info
	
	The following example shows an alert when a fabric device app state health is degraded:
	`<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/Device/Configuration/AppState?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31708"`
	`cause="configReady"`
	`type="fabricService"`
	`severity="major"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info="{AppState: cfg ready}"`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has app state set to cfg ready.`

Health Response	Response
	`{`
	
	`Resource: /App/Component/Fabric/Device/Configuration/AppState?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Red`
	`Value: 4`
	`}`
	`StatusText: The device "10.x.x.x" of fabric "fb" has app state set to cfg ready`
	`}`

Table 30: Managed Fabric Device Appstate Healthy Notice

31709	Managed Fabric Device Appstate Healthy Notice
Description	Send an alert when fabric device app state health is changed from Red or Black to Green.
Preconditions	Fabric is created, and devices are added and configured in XCO so that the App state changes to `cfg-in-sync`. The alert has the following app states and the corresponding severity: **Table 30: Managed Fabric Device Appstate Healthy Notice** State: cfg-in-sync — Severity: Info — Description: Device configurations are pushed to switch and it is in sync with XCO .

Requirements	Alert shows the following data:
	• Fabric Name
	• Device IP
	• Fabric Health Info
	The following example shows an alert when a fabric device app state is healthy:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/Device/Configuration/AppState?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31709"`
	`cause="appStateHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info=""`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has app state set to healthy.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Device/Configuration/AppState?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The device "10.x.x.x" of fabric "fb" has app state set to healthy`
	`}`

Table 31: Managed Fabric Device Provisioning State Degraded Notice

31710	Managed Fabric Device Provisioning State Degraded Notice
Description	Send an alert when fabric device provisioning state health is changed from Green to Black or Red.
Preconditions	Fabric is created, and devices are added in XCO. The alert has the following app states and the corresponding severity: **Table 31: Managed Fabric Device Provisioning State Degraded Notice** State: Not Provisioned — Severity: Major State: Provisioning Failed — Severity: Critical
Requirements	Alert shows the following data: • Fabric Name • Device IP • Fabric Health Info The following example shows an alert when a fabric device provisioning state health is degraded: `<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric/Device/Configuration/DevState?fabric_name=fb&device_ip=10.x.x.x" alertId="31710" cause="notProvisioned" type="fabricService" severity="major"] [alertData@1916 fabric_name="fb" device_ip="10.x.x.x" fabric_health_info="{ DevState: not provisioned}"] BOM The device "10.x.x.x" of fabric "fb" has dev state set to not provisioned.`

Health Response	Response
	`{`
	
	`Resource: /App/Component/Fabric/Device/Configuration/DevState?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Red`
	`Value: 4`
	`}`
	`StatusText: The device "10.x.x.x" of fabric "fb" has dev state set to not provisioned.`
	`}`

Table 32: Managed Fabric Device Provisioning State Healthy Notice

31711	Managed Fabric Device Provisioning State Healthy Notice
Description	Send an alert when fabric device provisioning state health is changed from Black or Red to Green.
Preconditions	Fabric is created, and devices are added and configured in XCO so that the devices move to provisioned state **(Severity - Info)**.

Requirements	Alert shows the following data:
	• Fabric Name
	• Device IP
	• Fabric Health Info
	
	The following example shows an alert when a fabric device provisioning state is healthy:
	`<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="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/Device/Configuration/DevState?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31711"`
	`cause="devStateHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info=""`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has dev state set to healthy.`

Health Response	Response
	`{`
	
	`Resource: /App/Component/Fabric/Device/Configuration/DevState?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The device "10.x.x.x" of fabric "fb" has dev state set to healthy.`
	`}`

Managed Fabric Device MCT Cluster Degraded Notice

31712	Managed Fabric Device MCT Cluster Degraded Notice
Description	Send an alert when fabric device MCT cluster health is changed from Green to Red or Black.
Preconditions	Fabric is created and MCT devices are added and configured in XCO. An alert is raised if any of the following cluster operational states are down: • PeerState: false • PeerKeepAliveState: false • ClusterState: false

Peer State	Cluster State	Peer Keep-alive State	Severity
Up	Up	Up	Info (Raised by alert ID 31519)
Up	Up	Down	Major
Up	Down	Up	Major
Up	Down	Down	Major
Down	Up	Up	Critical
Down	Up	Down	Critical
Down	Down	Up	Critical
Down	Down	Down	Critical

31712	Managed Fabric Device MCT Cluster Degraded Notice

Requirements	Alert shows the following data:
	1. Fabric Name
	2. Device IP
	3. Fabric Health Info
	The following example shows an alert when a fabric device MCT cluster state is degraded:
	`<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916 resource="/App/Component/Fabric/Device/Operational/Mct?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31712"`
	`cause="mctClusterStateAndPeerKeepAliveStateDown"`
	`type="fabricService"`
	`severity="major"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info="{`
	`PeerKeepAliveState: false`
	`ClusterState: false`
	`}"`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has MCT cluster state and peer keep alive state down.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Device/Operational/Mct?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Red`
	`Value: 4`
	`}`
	`StatusText The device "10.x.x.x" of fabric "fb" has MCT cluster state and peer keep alive state down`
	`}`

Managed Fabric Device MCT Cluster Healthy Notice

31713	Managed Fabric Device MCT Cluster Healthy Notice
Description	Send an alert when fabric device MCT cluster health is changed from Red or Black to Green.
Preconditions	Fabric is created and MCT devices are added and configured in XCO. An alert is raised if all the cluster operational states are up (**Severity: Info**): • PeerState: true • PeerKeepAliveState: true • ClusterState: true

31713	Managed Fabric Device MCT Cluster Healthy Notice

Requirements	Alert shows the following data:
	1. Fabric Name
	2. Device IP
	3. Fabric Health Info
	The following example shows an alert when a fabric device MCT cluster state is healthy:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/Device/Operational/Mct?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31713"`
	`cause="deviceMctClusterHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info=""`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has MCT cluster in healthy state.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Device/Operational/Mct?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText The device "10.x.x.x" of fabric "fb" has MCT cluster in healthy state`
	`}`

Managed Fabric Device Physical Topology Degraded Notice

31714	Managed Fabric Device Physical Topology Degraded Notice
Description	Send an alert when fabric device physical topology health is changed from Green to Red.
Preconditions	Fabric is created and devices are added and configured in XCO. Alerts are raised if the physical topology validation fails. The **severity** for physical topology errors is **Major**. Alerts are raised if the following conditions are not met: 1. Device level physical topology validations for non-Clos fabric: • Two devices in rack must have link between them. • Each rack must be connected to at least another rack. 2. 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 Super-spine nodes but not both. • Spine node must be connected to all the Border Leaf nodes. • More than two Leaf nodes must not be connected to each other. • 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. • Super Spine nodes must not be connected to each other. • If a Leaf node is "multi-homed", then the node must have an MCT neighbor. • If a Leaf node is "single-homed", then the node must not be connected to other Leaf nodes. • If a Border Leaf node is "multi-homed", then the node must have an MCT neighbor. • If a Border Leaf node is "single-homed", then the node must not be connected to other Border Leaf nodes.

31714	Managed Fabric Device Physical Topology Degraded Notice

Requirements	Alert shows the following data:
	1. Fabric Name
	2. Device IP
	3. Fabric Health Info
	The following example shows an alert when a fabric device physical topology health is degraded:
	`<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916 resource="/App/Component/Fabric/Device/Operational/Topology/Physical?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31714"`
	`cause="missingLinks"`
	`type="fabricService"`
	`severity="major"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info=":{`
	`Errors:[`
	`{`
	`Destination_ip: 10.x.x.a`
	`Destination_Role: leaf`
	`Device_links:[`
	`{Error: missing_links}`
	`]`
	`},`
	`Destination_ip: 10.x.x.b`
	`Destination_Role: leaf`
	`Device_links:[`
	`{Error: missing_links}`
	`]`
	`]`
	`"`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has missing-links with devices [10.x.x.a,10.x.x.b].`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Device/Operational/Topology/Physical?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Red`
	`Value: 4`
	`}`
	`StatusText The device "10.x.x.x" of fabric "fb" has missing-links with devices [10.x.x.a,10.x.x.b].`
	`}`

Managed Fabric Device Physical Healthy Notice

31715	Managed Fabric Device Physical Healthy Notice
Description	Send an alert when fabric device physical health is changed from Red to Green.
Preconditions	Fabric is created and devices are added and configured in XCO. Alerts are raised if the device physical topology validation is successful. **The severity for device physical topology errors is Info.** Alerts are raised if the following conditions are met: 1. Device level physical topology validations for non-Clos fabric: • Two devices in rack must have link between them. • Each rack must be connected to at least another rack. 2. 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 Super-spine nodes but not both. • Spine node must be connected to all the Border Leaf nodes. • More than two Leaf nodes must not be connected to each other. • 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. • Super Spine nodes must not be connected to each other. • If a Leaf node is "multi-homed", then the node must have an MCT neighbor. • If a Leaf node is "single-homed", then the node must not be connected to other Leaf nodes. • If a Border Leaf node is "multi-homed", then the node must have an MCT neighbor. • If a Border Leaf node is "single-homed", then the node must not be connected to other Border Leaf nodes.

31715	Managed Fabric Device Physical Healthy Notice

Requirements	Alert shows the following data:
	1. Fabric Name
	2. Device IP
	3. Fabric Health Info
	The following example shows an alert when a fabric device physical health is Green (healthy):
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/Device/Operational/Topology/Physical?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31715"`
	`cause="devicePhysicalTopologyHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info=""`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has physical topology in healthy state.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Device/Operational/Topology/Physical?fabric_name=fb&device_ip=10.x.x.x.a`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText The device "10.x.x.x" of fabric "fb" has physical topology in healthy state.`
	`}`

Managed Fabric Device Underlay Degraded Notice

31716	Managed Fabric Device Underlay Degraded Notice
Description	Send an alert when fabric device underlay health is changed from Green to Red or Black.
Preconditions	Fabric is created and devices are added and configured in XCO. The alerts are raised if the session state of any of the BGP neighbors is not established. • If BFP neighbors are not configured and devices are not in the provisioned state, then the severity is **Major**. • If some of the session state of BGP neighbor is down and some of the session state is up between devices, then the severity is **Major**. • If all the BGP neighbor sessions are down between devices, then the severity is **Critical**. • If the BGP neighbors are not configured and devices are in provisioned state, then the severity is **Critical**.

31716	Managed Fabric Device Underlay Degraded Notice
Requirements	Alert shows the following data: 1. Fabric Name 2. Device IP 3. Fabric Health Info The following example shows an alert when a fabric device underlay health is degraded: `<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric/Device/Operational/ Topology/Underlay?fabric_name=fb&device_ip=10.x.x.x" alertId="31716" cause="underlayNeighborsNotConfigured" type="fabricService" severity="major"] [alertData@1916 fabric_name="fb" device_ip="10.x.x.x" fabric_health_info=":{Errors:[{ Destination_ip: 10.x.x.a Destination_Role: leaf NeighborIP:10.x.x.2 SourceAsn: 6512, DeestinationAsn: 6500 Error: neighbor_not_configured },{ Destination_ip: 10.x.x.b Destination_Role: leaf NeighborIP:10.x,x.3 SourceAsn: 6512, DeestinationAsn: 6500 Error: neighbor_not_configured }],"] BOM The device "10.x.x.x" of fabric "fb" does not have bgp neighbors configured with [10.x.x.a,10.x.x.b].`
Health Response	Response `{ Resource: /App/Component/Fabric/Device/Operational/ Topology/Underlay?fabric_name=fb&device_ip=10.x.x.x HQI { Color: Red Value: 4 } StatusText The device "10.x.x.x" of fabric "fb" does not have bgp neighbors configured with [10.x.x.a,10.x.x.b] }`

Managed Fabric Device Underlay Healthy Notice

31717	Managed Fabric Device Underlay Healthy Notice
Description	Send an alert when fabric device underlay health is changed from Red or Black to Green.
Preconditions	Fabric is created and devices are added and configured in XCO. The alerts are raised if the BGP neighbors session state are in established state. The severity is **Info**.

Requirements	Alert shows the following data:
	1. Fabric Name
	2. Device IP
	3. Fabric Health Info
	The following example shows an alert when a fabric device underlay health is Green (healthy):
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId=”47”]`
	`[origin ip=”10.x.x.x” enterpriseId=”1916” software=”XCO” swVersion=“3.4.0”]`
	`[alert@1916`
	`resource=”/App/Component/Fabric/Device/Operational/Topology/Underlay?fabric_name=fb&device_ip=10.x.x.x”`
	`alertId=”31717”`
	`cause=” deviceUnderlayTopologyHealthRestored”`
	`type=“fabricService”`
	`severity=”info”]`
	`[alertData@1916`
	`fabric_name=”fb”`
	`device_ip=”10.x.x.x”`
	`fabric_health_info=””`
	`]`
	`BOM The device 10.x.x.x of fabric fb has underlay topology healthy.`

Health Response	Response
	`Resource: /App/Component/Fabric/Device/Operational/Topology/Underlay?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText The device “10.x.x.x” of fabric “fb” has underlay topology in healthy state`
	`}`

Managed Fabric Device Overlay Degraded Notice

31718	Managed Fabric Device Overlay Degraded Notice
Description	Send an alert when fabric device overlay health is changed from Green to Black.
Preconditions	Fabric is created, devices are added and configured, and tenant L2 services are configured with common ctag range in XCO. The alerts **(Severity: Critical)** are raised if the operational or admin status is down.

31718	Managed Fabric Device Overlay Degraded Notice
Requirements	Alert shows the following data: 1. Fabric Name 2. Device IP 3. Fabric Health Info The following example shows an alert when a fabric device overlay health is degraded: `<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="3.4.0"] [alert@1916 resource="/App/Component/Fabric/Device/Operational/Topology/Overlay?fabric_name=fb&device_ip=10.x.x.x" alertId="31718" cause="ovelayTunnelOperDown" type="fabricService" severity="critical"] [alertData@1916 fabric_name="fb" device_ip="10.x.x.x" fabric_health_info=":{ Errors:[{ Destination_ip: 10.x.x.a Destination_Role: leaf NeighborIP:10.10.10.2 SourceVTEPIP:10.10.10.3, DestinationVTEPIP:10.10.10.6 Admin_state:up, Oper_state: down Error: tunnel_oper_down },{ Destination_ip: 10.x.x.b Destination_Role: leaf SourceVTEPIP:10.x.x.3, DestinationVTEPIP:10.x.x.5 Admin_state:up, Oper_state: down Error: tunnel_oper_down }]"] BOM The device "10.x.x.x" of fabric "fb" does not have tunnels operationally up with [10.x.x.a,10.x.x.b]`
Health Response	Response `{ Resource: /App/Component/Fabric/Device/Operational/Topology/Overlay?fabric_name=fb&device_ip=10.x.x.x HQI { Color: Black Value: 5 } StatusText The device "10.x.x.x" of fabric "fb" does not have tunnels operationally up with [10.x.x.a,10.x.x.b] }`

Managed Fabric Device Overlay Healthy Notice

31719	Managed Fabric Device Overlay Healthy Notice
Description	Send an alert when fabric device overlay health is changed from Black to Green.
Preconditions	Fabric is created, devices are added and configured, and tenant L2 services are configured with common ctag range in XCO. The alerts **(Severity: Info)** are raised if the tunnel operational and admin status are up.

Requirements	Alert shows the following data:
	1. Fabric Name
	2. Device IP
	3. Fabric Health Info
	The following example shows an alert when a fabric device overlay health is Green (healthy):
	`<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="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/Device/Operational/Topology/Overlay?fabric_name=fb&device_ip=10.x.x.x"`
	`alertId="31719"`
	`cause=" deviceOverlayTopologyHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`device_ip="10.x.x.x"`
	`fabric_health_info=""`
	`]`
	`BOM The device "10.x.x.x" of fabric "fb" has overlay topology in healthy state.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/Device/Operational/Topology/Overlay?fabric_name=fb&device_ip=10.x.x.x`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText The device "10.x.x.x" of fabric "fb" has overlay topology in healthy state.`
	`}`

Managed Fabric Deleted Health Notice

31701	Managed Fabric Deleted Health Notice
Description	Send an alert when fabric is deleted.
Preconditions	Deletion of an existing fabric.

31701	Managed Fabric Deleted Health Notice

Requirements	Alert shows the following data:
	• Fabric Name
	The following example shows an alert when a fabric is deleted:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com AppFaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric?fabric_name=fb"`
	`alertId="31701"`
	`cause="configRemoved"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`]`
	`BOM : The fabric "fb" is deleted.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric?fabric_name=fb`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText : The fabric "fb" is deleted.`
	`}`

Managed Fabric Device Added Health Notice

31702	Managed Fabric Device Added Health Notice
Description	Send an alert when a fabric device is added.
Preconditions	Fabric is created and devices are added.

31702	Managed Fabric Device Added Health Notice
Requirements	Alert shows the following data: • Fabric Name • Device IP The following example shows an alert when a device is added to the fabric: `<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric/Device? fabric_name=fb&device_ip=10.x.x.x" alertId="31702" cause="configCreated" type="fabricService" severity="info"] [alertData@1916 fabric_name="fb" device_ip="10.x.x.x"] BOM The device "10.x.x.x" is added to the fabric "fb"`
Health Response	Response `{ Resource: /App/Component/Fabric/Device? fabric_name=fb&device_ip=10.x.x.x HQI { Color: Green Value: 0 } StatusText : The device "10.x.x.x" is added to the fabric "fb" }`

Managed Fabric Device Removed Health Notice

31703	Managed Fabric Device Removed Health Notice
Description	Send an alert when a fabric device is deleted.
Preconditions	Fabric is created. Devices are added and then deleted.

31703	Managed Fabric Device Removed Health Notice
Requirements	Alert shows the following data: • Fabric Name • Device IP The following example shows an alert when an existing fabric is deleted: `<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric/Device?fabric_name=fb&device_ip=10.x.x.x" alertId="31703" cause="configRemoved" type="fabricService" severity="info"] [alertData@1916 fabric_name="fb" device_ip="10.x.x.x"] BOM The device "10.x.x.x" is removed from the fabric "fb"`
Health Response	Response `{ Resource: /App/Component/Fabric/Device?fabric_name=fb&device_ip=10.x.x.x HQI { Color: Green Value: 0 } StatusText : The device "10.x.x.x" is removed from fabric "fb". }`

Fabric State Degraded Notice

31704	Fabric State Degraded Notice
Description	Send an alert when a fabric state health is changed from Green to Red or Black.
Preconditions	Fabric is created and devices are added in XCO. The following are the states and the severities of an alert: • Configure-failed: Fabric configure fails **(Critical)** • Migrate-success: fabric is migrated **(Major)** • Migrate-failed: fabric migrate failed **(Major)** • Settings-updated: fabric settings is updated **(Major)**

31704	Fabric State Degraded Notice
Requirements	Alert shows the following data: • Fabric Name • Fabric Health Info The following example shows an alert when an existing fabric is deleted: `<114> 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric/State?fabric_name=fb" alertId="31704" cause="fabricSettingsUpdated" type="fabricService" severity="major"] [alertData@1916 fabric_name="fb" fabric_health_info="{Fabric_status: setting updated}"] BOM The fabric "fb" has status set to settings-updated due to change in fabric settings [BGP-MD5].`
Health Response	Response `{ Resource: /App/Component/Fabric/State?fabric_name=fb HQI { Color: Red Value: 4 } StatusText: The fabric "fb" has status set to settings-updated due to change in fabric settings [BGP-MD5] }`

Managed Fabric State Healthy Notice

31705	Managed Fabric State Healthy Notice
Description	Send an alert when fabric state health is changed from Red or Black to Green.
Preconditions	Fabric is created and devices are added in XCO. The alert has the following state changes: • Configure-success: Fabric configure success **(Info)**

31705	Managed Fabric State Healthy Notice

Requirements	Alert shows the following data:
	• Fabric Name
	• Fabric Health Info
	The following example shows an alert when a fabric is created:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric/State?fabric_name=fb"`
	`alertId="31705"`
	`cause="fabricStateHealthRestored"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`fabric_health_info=""`
	`]`
	`BOM The fabric "fb" has status set to healthy.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric/State?fabric_name=fb`
	`HQI {`
	`Color:Green`
	`Value: 0`
	`}`
	`StatusText:The fabric "fb" has status set to healthy`
	`}`

Managed Fabric Created Notice

31700	Managed Fabric Created Notice
Description	Send an alert when a fabric is created.
Preconditions	None

31700	Managed Fabric Created Notice

Requirements	Alert shows the following data:
	• Fabric Name
	The following example shows an alert when a fabric is created:
	`<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -`
	`[meta sequenceId="47"]`
	`[origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"]`
	`[alert@1916`
	`resource="/App/Component/Fabric?fabric_name=fb"`
	`alertId="31700"`
	`cause="configCreated"`
	`type="fabricService"`
	`severity="info"]`
	`[alertData@1916`
	`fabric_name="fb"`
	`]`
	`BOM : The fabric "fb" is created.`

Health Response	Response
	`{`
	`Resource: /App/Component/Fabric?fabric_name=fb`
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText : The fabric "fb" is created.`
	`}`

Managed Fabric Health Degraded Notice

31799	Managed Fabric Health Degraded Notice
Description	Send an alert there is a change in fabric health or its contributors.
Preconditions	Fabric is created and devices are added and configured in the fabric. Severity is based on the following fabric health color value: • Red: **Major** • Black: **Critical** • Green: **Info** (Raised by the alert ID 31800)

31799	Managed Fabric Health Degraded Notice

Requirements	Alert shows the following data:
	• Fabric Name
	• Fabric Health Info
	The following example shows an alert when a fabric health is degraded:
	<114>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - -
	[meta sequenceId=”47”]
	[origin ip=”10.x.x.x” enterpriseId=”1916” software=”XCO” swVersion=“3.4.0”]
	[alert@1916
	resource=”/App/Component/Fabric?fabric_name=fb”
	alertId=”31799”
	cause=”fabricHealthCritical”
	type=“fabricService”
	severity=”critical”]
	[alertData@1916
	fabric_name=”fb”
	fabric_health_info=” Name : fb
	Health: Black
	Devices:[
	{
	Ip: 10.x.x.1,
	Device_health: Black
	Oper_state_health:{
	Underlay_health:{
	Health: Black
	Errors:[
	Source_ip: 10.x.x.1
	Source_asn: 6500
	Destination_asn: 6512
	Destination_ip: 10.x.x.2
	Neighbor_ip: 10.x.x.y
	Underlay_state: CONN
	Error: session_not_established
	
	]
	},
	
	}
	}
	
	]
	”
	]
	BOM : Fabric "fb" health is in critical state because the underlay topology has errors for the device(s) [10.x.x.1,10.x.x.2]

Health Response	Response
	{
	Resource: /App/Component/Fabric?fabric_name=fb
	HQI {
	Color: Black
	Value: 5
	}
	StatusText : Fabric "fb" health is in critical state because the underlay topology has errors for the device(s) [10.x.x.1,10.x.x.2].

Managed Fabric Health Restored Alert Notice

31800	Managed Fabric Health Restored Alert Notice
Description	Send an alert when the fabric health is changed to Green.
Preconditions	Fabric is created and devices are added and configured in the fabric.
Requirements	Alert shows the following data: • Fabric Name • Fabric Health Info The following example shows an alert when a fabric health is restored: `<118>1 2003-10-11T22:14:15.003Z xco.machine.com FaultManager - - [meta sequenceId="47"] [origin ip="10.x.x.x" enterpriseId="1916" software="XCO" swVersion="3.4.0"] [alert@1916 resource="/App/Component/Fabric?fabric_name=fb" alertId="31800" cause="fabricHealthRestored" type="fabricService" severity="info"] [alertData@1916 fabric_name="fb" fabric_health_info=""] BOM : The fabric "fb" health is restored.`
Health Response	Response `{ Resource: /App/Component/Fabric?fabric_name=fb HQI { Color: Green Value: 0 } StatusText : The fabric "fb" health is restored.`

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="3.4.0"] [alert@1916 resource="/App/System/HA/Nodes/Node" alertId="31050" cause="lossOfRedundancy" type="operationalViolation" severity="minor"] [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: Orange
	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="3.4.0"]`
	`[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=”3.4.0”] [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: 4
	}
	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=”3.4.0”] [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: 2
	}
	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="3.4.0"] [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.
	}

License Alerts

Use the information in the following tables to learn about all possible license alerts in detail that are raised by Fault Management.

License Expiry Threshold

31400	License Expiry Threshold
Description	Send an alert per AID when expiry is 90, 60, 30 or less days away.
Preconditions	• Polling frequency for license expiry threshold notice: **daily** • License Expiry Thresholds: 1. **90 days - warning** 2. **60 days - minor** 3. **30 or less days - major** The daily polling sends an "Alert" event notification with the license AID which is processed by the fault engine.

31400	License Expiry Threshold
Requirements	Alert shows the following data: • Expiry Date • AID • License Filename • License Feature The following example shows an alert when a login attempt to XCO fails: `<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="3.5.0"] [alert@1916 resource="/App/System/Security/License?aid=5272f34d-e1fd-488c-bf6c-7e2b2d831db1 alertId="31400" cause="licenseExpiryThreshold" type=" communicationsAlarm" severity="warning"] [alertData@1916 aid="5272f34d-e1fd-488c-bf6c-7e2b2d831db1" expiry_date="11 OCT 2024 23 59" license_filename="lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic" license_feature=" PRD-XCO-EFA-D1"] BOMThe license file lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic with feature type PRD-XCO-EFA-D1 and AID 5272f34d-e1fd-488c-bf6c-7e2b2d831db1 is 90 days from expiry date 11 OCT 2024 23 59`

Health Response	Response
	`{ Resource: /App/System/Security/License?aid=5272f34d-e1fd-488c-bf6c-7e2b2d831db1`
	
	`HQI { Color: Yellow Value: 2 } StatusText: The license with AID 5272f34d-e1fd-488c-bf6c-7e2b2d831db1 will expire in 90 days. }`
	
	Note: This alert does not impact XCO health in XCO 3.5.0.

License Expired

31401	License Expired
Description	Send an alert per AID when a license is expired.
Preconditions	The daily polling sends an "Alert" event notification with the license AID which is processed by the fault engine. Once the licenses are expired, this alert is sent every day until you renew the licenses. **Severity: Critical**

31401	License Expired
Requirements	Alert shows the following data: • Expiry Date • AID • License Filename • License Feature The following example shows an alert when a license is expired: <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="3.5.0"] [alert@1916 resource="/App/System/Security/License?**aid=5272f34d-e1fd-488c-bf6c-7e2b2d831db1** alertId="31400" cause="licenseExpiryThreshold" type=" communicationsAlarm" severity="warning"] [alertData@1916 aid="5272f34d-e1fd-488c-bf6c-7e2b2d831db1" expiry_date="11 OCT 2024 23 59" license_filename="lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic" license_feature=" PRD-XCO-EFA-D1"] BOMThe license file lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic with feature type PRD-XCO-EFA-D1 and AID 5272f34d-e1fd-488c-bf6c-7e2b2d831db1 has expired

Health Response	Response
	{
	Resource: /App/System/Security/License?**aid=5272f34d-e1fd-488c-bf6c-7e2b2d831db1**
	HQI {
	Color: Black
	Value: 5
	}
	StatusText: The license with AID 5272f34d-e1fd-488c-bf6c-7e2b2d831db1 has expired.
	}
	
	Note: This alert does not impact XCO health in XCO 3.5.0.

License Expiry Clear

31402	License Expiry Clear
Description	Send an alert per AID when a license is renewed and license expiry is more than 90 days away.
Preconditions	The daily polling sends an "Alert" event notification with the license AID which is processed by the fault engine. This will be sent to clear the LicenseExpiredAlert when you renew or delete a license. **Severity: Clear**

31402	License Expiry Clear

Requirements	Alert shows the following data:
	• Expiry Date
	• AID
	• License Name
	• License Feature
	
	The following example shows an alert when a license is renewed:
	`<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="3.5.0"]`
	`[alert@1916`
	`resource="/App/System/Security/License?`**`aid=5272f34d-e1fd-488c-bf6c-7e2b2d831db1`**
	`alertId="31400"`
	`cause="licenseExpiryThreshold"`
	`type=" communicationsAlarm"`
	`severity="warning"]`
	`[alertData@1916`
	`aid="5272f34d-e1fd-488c-bf6c-7e2b2d831db1"`
	`expiry_date="11 OCT 2024 23 59"`
	
	`license_filename="lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic"`
	`license_feature=" PRD-XCO-EFA-D1" ]`
	`BOMThe license file lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic with feature type PRD-XCO-EFA-D1 and AID 5272f34d-e1fd-488c-bf6c-7e2b2d831db1 has been cleared.`

Health Response	Response
	`{`
	`Resource: /App/System/Security/License?`**`aid=5272f34d-e1fd-488c-bf6c-7e2b2d831db1`**
	`HQI {`
	`Color: Green`
	`Value: 0`
	`}`
	`StatusText: The license with AID 5272f34d-e1fd-488c-bf6c-7e2b2d831db1 has been renewed or deleted.`
	`}`
	
	Note: This alert does not impact XCO health in XCO 3.5.0.

Password Expiry Alerts

Use the information in the following tables to learn about all possible password expiry alerts in detail that are raised by Fault Management.

Password Expiry Threshold

35100	Password Expiry Threshold
Description	Send an alert when an SLX user's password is about to expire.
Preconditions	GoRaslog receives the raslog message from SLX that indicates that the password for a given user it about to expire. The following alert levels are from payload of raslog. These are not severities of the alerts that will be raised. • info – if a password is about to expire before the number of days configured for Info alert level for a given user • minor – if a password is about to expire before the number of days configured for minor alert level for a given user • major – if a password is about to expire before the number of days configured for major alert level for a given user • critical – if a password is about to expire before the number of days configured for critical alert for a given user **Severity: Warning**

35100	Password Expiry Threshold
Requirements	Alert shows the following data: • Device IP • Username • Alert level • Days before expiry The following example shows an alert when a password is about to expire: `<116>1 2003-10-11T22:14:15.003Z xco.machine.com AppFaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.5.0"] [alert@1916 resource=/App/Component/Asset/Device/Password?device_ip=10.2.3.4&user_name=test1" alertId="35100" cause="passwordExpiryProblem" type="processingErrorAlarm" severity="warning"] [alertData@1916 user_name="test1" device_ip="10.2.3.4" alert_level="major" expiry_days=2 xco_device_user=true] BOMPassword for User "test1" will expire in 2 days on device 10.2.3.4.`
Health Response	Response `{ Resource: /App/Component/Asset/Device/Password?device_ip=10.2.3.4&user_name=test1 HQI { Color: Yellow Value: 2 } StatusText: Password for User "test1" will expire in 2 days on device 10.2.3.4.. }`

Password Expired

35101	Password Expired
Description	Send an alert when an SLX user's password has expired.
Preconditions	GoRaslog receives the raslog message from SLX that indicates that the password for a given user has expired. 0 days indicate that the password expired today. **Severity: Minor**

35101	Password Expired
Requirements	Alert shows the following data: • Device • Username • Days after Expiry The following example shows an alert when a password is expired: <116>1 2003-10-11T22:14:15.003Z xco.machine.com AppFaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.5.0"] [alert@1916 resource=/App/Component/Asset/Device/Password? device_ip=10.2.3.4&user_name=test1" alertId="35101" cause="keyExpired" type="processingErrorAlarm" severity="minor"] [alertData@1916 user_name="test1" device_ip="10.2.3.4" expired_days=2 xco_device_user=false] BOMPassword for User "test1" had expired 2 days ago on device 10.2.3.4.

Health Response	Response
	{
	Resource: /App/Component/Asset/Device/Password?
	device_ip=10.2.3.4&user_name=test1
	HQI {
	Color: Orange
	Value: 3
	}
	StatusText: Password for User "test1"
	had expired 2 days ago on device 10.2.3.4...
	}

Password Expiry Clear

35102	Password Expiry Clear
Description	Send an alert when an SLX user's password is renewed.
Preconditions	GoRaslog receives the raslog message from SLX that indicates that the password for a given user has been renewed. **Severity: Info**

35102	Password Expiry Clear
Requirements	Alert shows the following data: • Device IP • Username The following example shows an alert when a password is renewed: <116>1 2003-10-11T22:14:15.003Z xco.machine.com AppFaultManager - - [meta sequenceId="47"] [origin ip="10.20.30.40" enterpriseId="1916" software="XCO" swVersion="3.5.0"] [alert@1916 resource=/App/Component/Asset/Device/Password? device_ip=10.2.3.4&user_name=test1" alertId="35102" cause="passwordExpiryClear" type="processingErrorAlarm" severity="info"] [alertData@1916 user_name="test1" device_ip="10.2.3.4" xco_device_user=true]] BOMPassword for User "test1" renewed on device 10.2.3.4.
Health Response	Response { Resource: /App/Component/Asset/Device/Password? device_ip=10.2.3.4&user_name=test1 HQI { Color: Green Value: 0 } StatusText: Password for User "test1" renewed on device 10.2.3.4 }

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="3.4.0"] [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: 4
	}
	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="3.4.0"] [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

31033	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.

31033	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="3.4.0"] [alert@1916 resource="/App/System/Security/Authentication?server=10.1.2.3" alertId="31033" 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. }

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=”3.4.0”] [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=”3.4.0”] [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

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="3.4.0"] [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: 2 } 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="3.4.0"] [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: 5 } 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=”3.4.0”] [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=”3.4.0”] [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=”3.4.0-v1” new_version=”3.4.0-v2”] BOMDeployment Suite "fabric" upgrade initiated for type "multi-node" on Platform "TPVM" with Original Version "3.4.0-v1" to New Version "3.4.0-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=”3.4.0”]`
	`[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=”3.4.0-v1”`
	`new_version=”3.4.0-v2”]`
	`BOMDeployment Suite "fabric" upgrade successfully completed for type "multi-node" on Platform "TPVM" with Original Version "3.4.0-v1" to New Version "3.4.0-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

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.

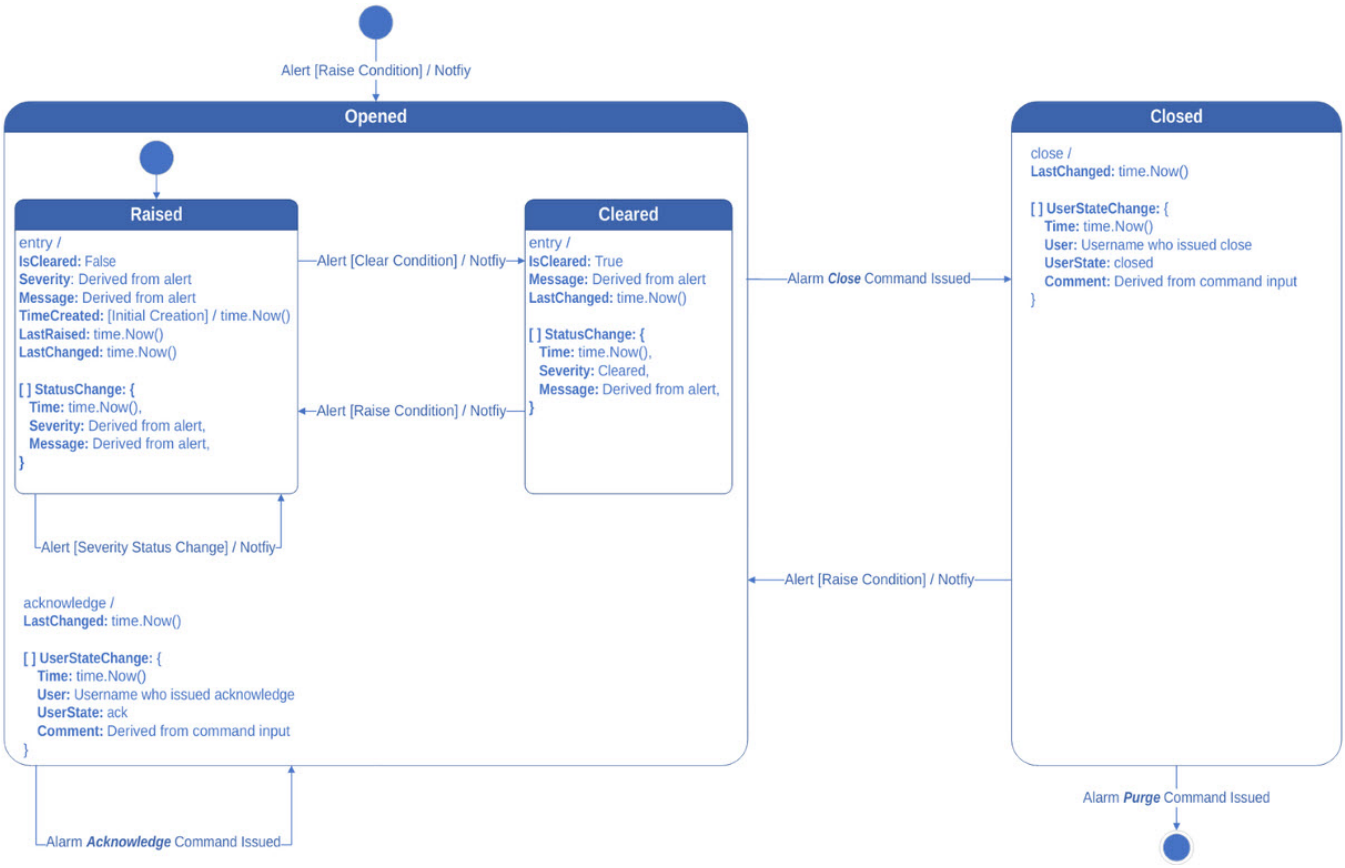

Figure 57: XCO alarm state transitions

Note

XCO supports a maximum of 500 alarms on TPVM deployments. You must purge the closed alarms manually to allow new alarms.

Alarm Severity

The following tables describe the severity levels of an alarm:

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:** When assigning a severity to an alert, do not use this value.
Cleared	1	Alarm is cleared by the system. **Note:** This severity level is only applicable to the "StatusChange" list severity.

Alarm Types

The following tables describe the types of alarm:

Types	Enum	Description
Policy	14	Indicates a policy-service related issue
Tenant	13	Indicates a tenant-service related issue
Fabric	12	Indicates a fabric-service related issue
TimeDomain	11	Indicates that an event has occurred at an unexpected or prohibited time
Security	10	Indicates a security violation such as authentication failure or unauthorized access attempt
Physical	9	Indicates cable tampering or intrusion
Operational	8	Indicates that the provisioning of the requested service was not possible due to unavailability or malfunction of the service
Integrity	7	Indicates duplicate, missing, modified, unexpected, or out of sequence information
Environmental	6	Indicates an issue related to the enclosure housing the equipment
Equipment	5	Indicates faulty equipment

Types	Enum	Description
Processing	4	Indicates a software processing issue
QualityOfService	3	Indicates a quality related issue
Communications	2	Indicates a communication related issue
Other	1	Indicates a catch-all category for alarms.

Alarm Inventory

The alarm inventory contains a list of system default alarms.

Certificate Expiration

ID	32000
Type	Security
Max Alarm Instance	7 (application certificate types)
Description	Raise an alarm to notify that an XCO certificate is about to expire or has expired.
Associated Alerts	**XCO 3.2.0:** CertificateExpiryNoticeAlert CertificateExpiredAlert CertificateRenewalAlert **XCO 3.3.0:** CertificateUnreadableAlert
Severity	Warning / Critical
Will Clear	True
Raise or Clear Conditions (Status Change)	**Warning:** Expiration of an XCO certificate within 90 days (Certificate Expiry Notice Alert) or Error reading the certificate during component polling (Certificate Unreadable Alert) **Critical:** XCO certificate has expired (Certificate Expired Alert) **Cleared:** XCO certificate renewal (Certificate Renewal Alert)
Action	Notify APP_ALARMS syslog / webhook

Device Certificate Expiration

ID	32001
Type	Security
Max Alarm Instance	Number of Devices * 3 (device certificate types)
Description	Raise an alarm to notify that a device certificate is about to expire or has expired.

Associated Alerts	**XCO 3.2.0:** DeviceCertificateExpiryNoticeAlert DeviceCertificateExpiredAlert DeviceCertificateRenewalAlert **XCO 3.3.0:** DeviceCertificateUnreadableAlert DeviceCertificateDeviceRemovedAlert
Severity	Warning / Critical
Will Clear	True
Raise or Clear Conditions (Status Change)	
Action	Notify APP_ALARMS syslog / webhook

Login Authentication

ID	32010
Type	Security
Max Alarm Instance	Number of Users* *This Includes existing and non-existing application users.
Description	Raise an alarm to notify suspicious login activity.
Associated Alerts	LoginFailureAlert
Severity	Warning
Will Clear	False
Raise or Clear Conditions (Status Change)	**Warning:** 5 successive failed login attempts within 1 minute **System will not clear the alarm.**
Action	Notify APP_ALARMS syslog / webhook

LDAP Server Connectivity

ID	32030
Type	Communication
Max Alarm Instance	4 (number of supported LDAP servers)
Description	Raise an alarm and notify an LDAP server is no longer reachable.
Associated Alerts	LDAPServerConnectivityFailureAlert LDAPServerConnectivitySuccessAlert LDAPServerConfigurationRemovedAlert
Severity	Major
Will Clear	True
Raise or Clear Conditions (Status Change)	**Major:** LDAP server connectivity failure. **Cleared:** LDAP server connectivity restored or LDAP server configuration removal.
Action	Notify APP_ALARMS syslog / webhook

Table 33: Password Expiration

ID	35000
Type	Communications
Max Alarm Instance	1 (Per Device, SLX user combination)
Description	Raise an alarm when the user password is about to expire or already expired on a given device
Associated Alerts	DevicePasswordExpiryThresholdAlert DevicePasswordExpiredAlert DevicePasswordExpiryClearAlert
Severity	Warning/Minor
Will Clear	True
Raise or Clear Conditions (Status Change)	**Info:** When password is about to expire and have info alert level **Minor:** When password is about to expire and have minor alert level **Major:** When password is about to expire and have major alert level **Critical:** When password is about to expire and have major critical level **Cleared:** Password Expiry Clear
Action	Notify APP_ALARMS syslog or webhook

Storage Utilization

ID	32040
Type	Processing
Max Alarm Instance	Number of Nodes * 2 (monitored mount points)
Description	Raise an alarm to notify storage utilization for XCO installation has reached a certain threshold or it is full.
Associated Alerts	StorageUtilizationThresholdAlert StorageUtilizationFullAlert StorageUtilizationCheckAlert
Severity	Warning - Critical (All severities between Warning and Critical are possible)
Will Clear	True

Raise or Clear Conditions (Status Change)	**Warning:** File system utilization is 75% - 84% (Storage Utilization Threshold Alert) **Minor:** File system utilization is 85% - 94% (Storage Utilization Threshold Alert) **Major:** File system utilization is 95% - 96% (Storage Utilization Threshold Alert) **Critical:** File system utilization is 97% - 100% (Storage Utilization Threshold Alert or Storage Utilization Full Alert) **Cleared:** File system utilization is below 75% (Storage Utilization Check Alert)
Action	Notify APP_ALARMS syslog / webhook

Device Connectivity

ID	32500
Type	Communications
Max Alarm Instance	Number of Devices
Description	Raise an alarm and notify a registered device has lost connectivity either due to the configured device health setting or during a device inventory update.
Associated Alerts	DeviceConnectivityFailureAlert DeviceConnectivitySuccessAlert DeviceConnectivityDeviceRemovedAlert
Severity	Major
Will Clear	True
Raise or Clear Conditions (Status Change)	**Major:** Device connectivity failure. **Cleared:** Device connectivity restored or Device registration removal.
Action	Notify APP_ALARMS syslog / webhook

HA Service Redundancy

ID	32050
Type	Operational
Max Alarm Instance	1 (Single HA service alarm)
Description	Raise an alarm and notify a loss of redundancy and/or failover has occurred.
Associated Alerts	HAServiceNonRedundantAlert HAServiceFullyRedundantAlert HAServiceNewActiveAlert
Severity	Major
Will Clear	True
Raise or Clear Conditions (Status Change)	**Major:** HA redundancy is lost **Cleared:** HA redundancy is restored
Action	Notify APP_ALARMS syslog / webhook

Node Service

ID	32051
Type	Operational
Max Alarm Instance	1 (Single node service alarm)
Description	Raise an alarm and notify the nodes' services are not all running and degraded.
Associated Alerts	ServiceDegradedAlert ServiceRestoredAlert
Severity	Warning
Will Clear	True
Raise or Clear Conditions (Status Change)	**Major:** Nodes' services are not all running. **Cleared:** Nodes' services are all running
Action	Notify APP_ALARMS syslog / webhook

Fabric Health

ID	33000
Type	FabricService
Max Alarm Instance	Number of fabrics
Description	Raise an alarm and notify that fabric health is changed. **Note:** The fabric alarm and fabric alarm status update notifications indicate that the fabric alarm is cleared when it should remain raised, but quickly gets updated to the proper raised state.

Associated Alerts	FabricStateDegradedAlert FabricStateHealthyAlert FabricPhysicalTopologyDegradedAlert FabricPhysicalTopologyHealthyAlert FabricDeviceAppStateDegradedAlert FabricDeviceAppStateHealthyAlert FabricDeviceProvisioningStateDegradedAlert FabricDeviceProvisioningStateHealthyAlert FabricDeviceMctDegradedAlert FabricDeviceMctHealthyAlert FabricDevicePhysicalTopologyDegradedAlert FabricDeviceUnderlayTopologyDegradedAlert FabricDeviceOverlayTopologyDegradedAlert FabricDevicePhysicalTopologyHealthyAlert FabricDeviceUnderlayTopologyHealthyAlert FabricDeviceOverlayTopologyHealthyAlert FabricHealthDegradedAlert FabricHealthRestoredAlert FabricDeviceRemovedAlert FabricDeletedAlert
Severity	Major
Will Clear	True
Raise or Clear Conditions (Status Change)	**Major:** Fabric health Degraded **Cleared:** Fabric Health Restored
Action	Notify APP_ALARMS syslog / webhook

Port Flap

ID	34000
Type	Communications
Max Alarm Instance	1 (Per Device)
Description	1
Associated Alerts	PortFlapAlert PortFlapClearAlert
Severity	Major
Will Clear	True
Raise or Clear Conditions (Status Change)	**Major:** Port Flap Alert **Cleared:** Port Flap Clear Alert
Action	Notify APP_ALARMS syslog / webhook

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 port per device: 32

```
(CertificateExpiration Instances) + (DeviceCertificateExpiration Instances) +
(LoginAuthentication  Instances) + (StorageUtilization Instances) + (LDAPServer
Instances) + (DeviceConnectivity Instances) + (HAServiceRedunances Instance) +
(NodeService Instance)  + (PortFlap instance) =
(7) + (20 * 3) + (10) + (2 * 2) + (5) + (20) + (1) + (1) + (20 * 32)= 748
```

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
From XCO 3.3.0 onwards, system sends out similar notifications for APP_ALERTS and APP_ALARMS.

Table 34: 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
Minor	Minor	3	Error	Error condition
Warning	Warning	4	Warning	Warning condition
Indeterminate/ Cleared		5	Notice	Normal, but significant condition

Table 34: Syslog Severity (continued)

Alarm Severity	Alert Severity	Syslog Severity	Syslog Enum	Description
	Info	6	Informational	Informational messages
		7	Debug	Debug-level messages

Table 35: 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 #unique_590/unique_590_Connect_42_topic-870634b2-175a-4a40-adc0-6b5d55d7ace4 on page 1016
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
Software	origin	XCO	Software Name
SW Version	origin	3.5.0	Software Version
Resource	alarm	`/App/System/Security/Certificate?type=app_server_certificate`	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)

Table 35: Syslog Alarm (RFC-5674) - Common Alarm Payload (continued)

Field	SD-ID (Structured Data ID)	Example	Description
PerceivedSeverity	alarm	warning	Severity of the XCO Alarms See Alarm Severity on page 1009.
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 2023 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="3.5.0"]
[alarm resource="/App/System/Security/Certificate?type=app_server_certificate"
probableCause="keyExpired"
eventType="security"
perceivedSeverity="warning"]
[alarmData@1916
type="app_server_certificate"
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_certificate",
  "alarm_type": "security",
  "alarm_cause": "keyExpired",
  "alert_data": null,
  "alarm_data": {
    "type": "app_server_certificate",
    "expiry_date": "Sep 12 10:00:45 2022 GMT",
  }
}
```

Alarm Commands

XCO 3.4.0 supports the following system alarm commands:

Commands	Description
`efa system alarm inventory`	Lists supported alarms
`efa system alarm show`	Bulk show active alarms with resource parent or range of instance IDs
`efa system alarm summary`	Lists alarm counts and statistics
`efa system alarm acknowledge`	Bulk acknowledge the alarm instance with resource parent or range of instance IDs.
`efa system alarm close`	Bulk close the open alarms with resource parent or range of alarm instance IDs
`efa system alarm purge`	Purges all closed alarms

For more information, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

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.

Warning
A condition that triggered a Major alert in v4.0.2 may now appear as Minor in v4.1.0, because the updated engine accounts for MCT peer connectivity and redundancy paths. Alert delivery mechanisms and structures remain unchanged.

Fabric Health delivers a comprehensive, multi-layer view of Clos fabric status — covering physical topology, underlay routing, and overlay VXLAN tunnels. XCO 4.1.0 introduces enhancements that replace the previous three-level model (Healthy / Red / Black) with a five-level severity scale that more precisely reflects real network conditions. The enhancements provide the following business value:

- **Fewer false alarms** — redundancy-aware health evaluation means a single link or BGP failure no longer escalates to a Critical alert unnecessarily.
- **Early visibility into degradation** — Warning and Minor states surface partial failures before they affect service, giving teams a window to act.
- **Faster troubleshooting** — per-device, per-layer health reporting pinpoints the source of a problem quickly, reducing mean time to resolution.
- **Proactive remediation** — configurable link-count thresholds alert teams to weakened redundancy while traffic is still flowing, before an outage develops.

Bubbling of Health Status

In the following example,

- Certificate present in the `/App/System/Security/Certificate` directory is in Yellow state and goes up to the top level of the hierarchy.
- The `/App/System/HA/Nodes/Node` is Red.
- The `/App/System/Storage?mount_point=/&node_ip=10.20.247.101` is Yellow.
- The `/App/Component/Asset/Device?device_ip=10.20.246.30` is Red.
- The `/App/System` evaluates as Red since it is more critical between HA and Storage nodes.
- The `/App` evaluates as Red since both System and Component nodes are Red.

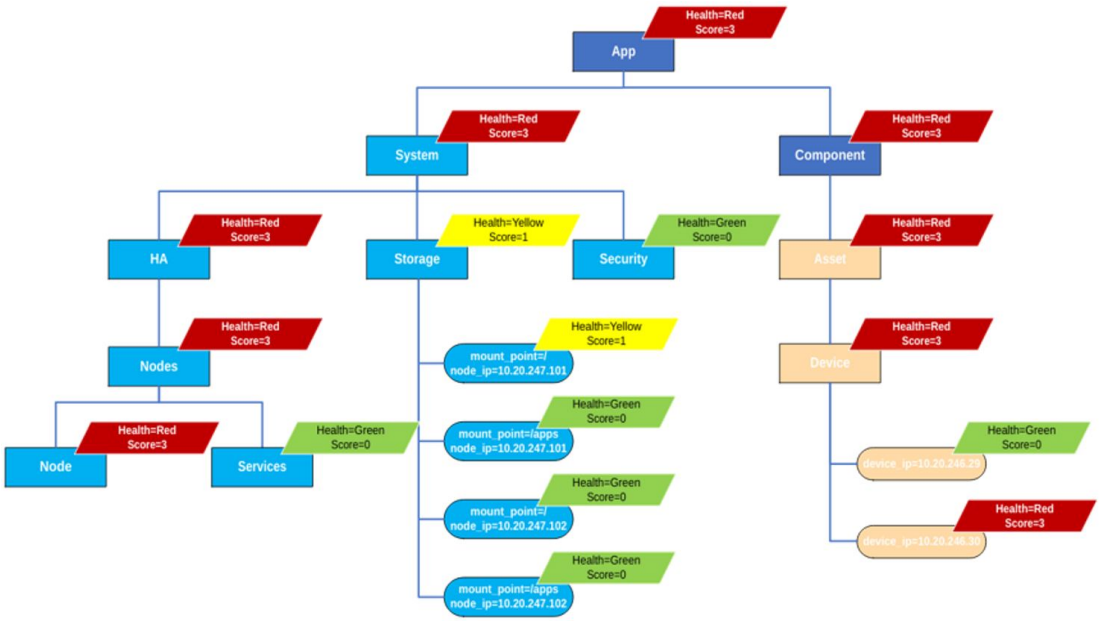

Each node in the hierarchy maintains the following details:

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. 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.

Table 36: Health Status Levels

HQI Color	HQI Value	Level	Description
Green	0	Healthy	Healthy System All physical, underlay, and overlay components are fully operational.
Gray	1		Unknown Value For example, fabric creation without adding devices.
Yellow	2	Warning	Potential Failure is imminent For example, certificate expiry in 10 days. Link counts have dropped below configured thresholds (min-p2p-link-count or min-mct-link-count). No traffic impact yet.
Orange	3	Minor	Failed (yet the system is functional) For example, standby node is down, partial degradation with no complete isolation. Traffic is still flowing, but redundancy is reduced.
Red	4	Major	Requires Immediate Attention For example, one or more devices are completely isolated from the fabric — full traffic loss on those nodes.
Black/Dark	5	Critical	All devices in the fabric (or within a single pod) are Major. Total loss of service.

Note

- XCO conducts regular polls to check the health of the system. If the health of XCO recovers within the time set for polling, then the health status remains Green. However, if any of the services are down at the time of polling, the XCO health status changes to "non-green". The health status turns green again if the services are up at the time of the next polling cycle. The current polling time is set to 60 seconds.
- When the active node's database is down, the **`efa health show`** command may not accurately reflect the XCO's health state.

Health Color Precedence Order

The overall XCO health status is determined by the highest-priority (most critical) health color reported across all monitored health indicators. When multiple resources report different health states, the **`efa health show`** command displays the color with the highest severity as the overall XCO health status.

The health color precedence order, from highest to lowest severity, is:

Critical (Black) > Degraded (Red) > Warning (Yellow) > Unknown/Indeterminate (Grey) > Healthy (Green)

For example, if the fabric health status is reported as **Grey** but a service, such as `/App/System/HA/Nodes/Services`, reports **Yellow**, the overall XCO health status is displayed as **Yellow**. This is because **Yellow** has a higher precedence than **Grey**. A Yellow status indicates a warning condition where one or more services are not fully operational, whereas Grey represents an unknown or indeterminate state. As a result, the warning condition is considered more significant and determines the overall health status.

Health Commands

XCO supports the following health commands/APIs:

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
`efa health debug clear --resource /App/Asset`	If a resource is not specified, the command clears the health of the tree from the Root node, for example, "/". If a resource is specified, the command clears the health of the subtree for a specified resource
`efa fabric setting update --min-p2p-link-count <N>`	Set P2P link Warning threshold
`efa fabric setting update --min-mct-link-count <N>`	Set MCT link Warning threshold
`efa fabric health calculate`	Manually trigger health recalculation
`GET /v1/fabric/fabric-health?name={name}`	Get health for a specific fabric
`GET /v1/fabric/fabrics-health`	Get health for all fabrics
`POST /v1/fabric/debug/health/compute`	Debug: force health recompute

Health APIs

Health service URI has additional query parameters to enable users to query for specific type of entity and asset. During API validation, the query parameters are validated and an appropriate response is returned. Check the REST API guide for the different status and responses.

Category	API	Description
Health – Configuration Information	**GET** ***/healthmanager/health/inventory***	Displays a list of contributor resources
Health – Operational Information	**GET** **/healthmanager/health/detail** **?** **resource=<resource path>**	Gets Health *Get Health Detail /App/ System/ Certificate* — Gets HQI values at certificate and any meta data providing details for health deterioration
Health – Operational Information	**GET** **/healthmanager/health** **?** **resource=<resource path>** **&** **detail=<true | false>**	Displays a list of resources with a health score contributing to overall health.
Health – Operational Information	**GET** **/v1/fabric/fabric-health?name={name}**	Health for a named fabric
Health – Configuration Information	**GET** **/v1/fabric/fabrics-health**	Health summary for all fabrics
Health – Configuration Information	**GET** **/v1/fabric/health**	General health endpoint
Health – Configuration Information	**POST** **/v1/fabric/debug/health/compute**	Trigger a debug health recompute

The following table describes commands to clear the health of a subtree for the specified resource:

Category	CLI	Description
Health – Configuration Information	Command Syntax: **efa health** debug clear -- resource /App/Asset	Optional –resource If a resource is not specified, the command clears the health of the tree from the Root node, for example, "/". If a resource is specified, the command clears the health of the subtree for a specified resource

Example

- Request Get Health Inventory

```
Response
{
    Resource: /App
    HQI {
        Color: Yellow
        Value: 2
    }
    StatusText: [<Freeform status text>]
}
```

- Request Get Health detail

```
Response
{
    Resource: /App
    HQI {
        Color: Yellow
        Value: 2
    }
   Contributor {
        ResourceList: [/App/System/Certificate, /App/Component/Fabric]
   }
   StatusText: [<Freeform status text>]
}
```

- Request Get Health resource=/App/System/Certificate

```
Response
{
    Resource:/App/System/Certificate
    HQI {
        Color: Yellow
        Value: 2
    }
    StatusText: [Certificate x expires on <date>.]
}
```

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 fabric health status is derived by aggregating device health states using the following precedence rules:

Fabric Status	Aggregation Rule
Critical	It takes the highest priority. All composed devices are Major, OR all devices within a single pod are Major.
Major	At least one device is Major (but not all).
Minor	At least one device is Minor; no device is Major.
Warning	At least one device is Warning; no device is Minor or Major.
Healthy	All devices are Healthy.

Note

- Green status indicates a healthy system.
- Red status indicates that the system is degraded and 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.

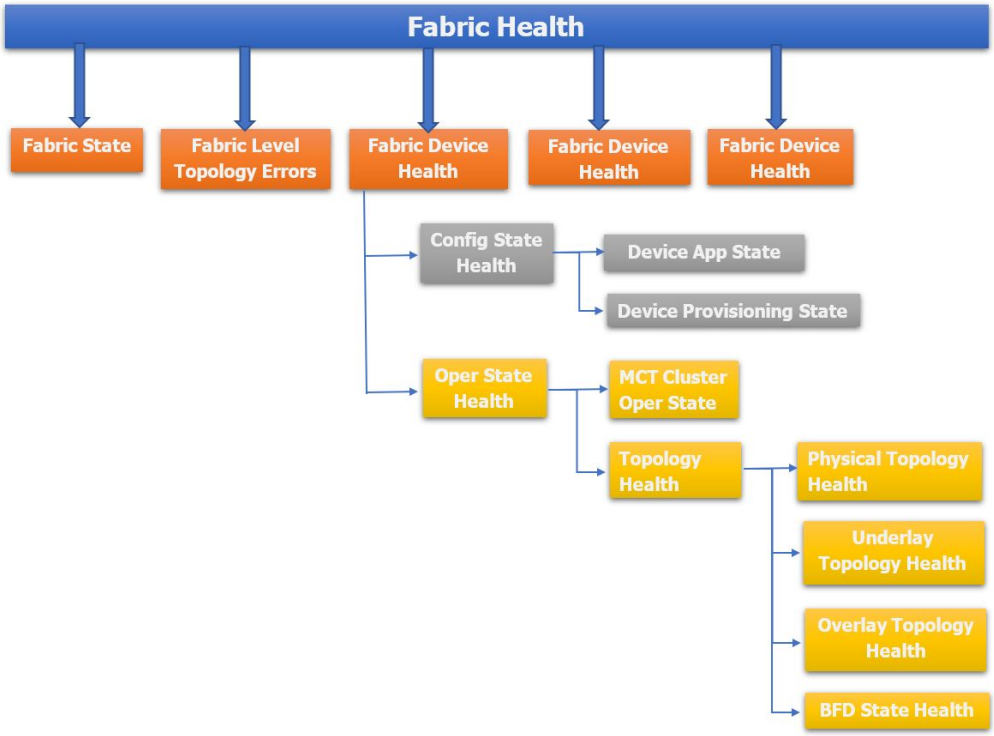

Figure 58: Fabric health hierarchy

The following is the truth table for fabric health:

Table 37: Truth Table

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
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

Table 37: Truth Table (continued)

Fabric Status	Fabric Level Physical Topology Health	Device Health	Fabric Health
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

Triggering a Health Recalculation

Health status is updated automatically on network events, but can also be triggered manually:

Trigger	Description
Event-driven	BGP session or link state changes trigger an immediate asynchronous recalculation.
Periodic timer	A background reconciliation loop corrects any missed events.
Manual CLI	Run the **`efa fabric health calculate`** command to force an immediate recalculation.

Fabric Status

The following table shows the mapping of a fabric status to the fabric health.

Fabric Status	Health
created (without devices added)	Green
created (with devices added)	Red
configure-success	Green
configure-failed	Black
settings-updated	Red
migrate-success	Red
migrate-failed	Red
upgrade-failed	Red
upgrade-success	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	

Service Health

Service health measures the availability of tenant services running over VXLAN tunnels within the fabric. It is assessed per service across all participating Leaf/Border-Leaf nodes:

Service Status	Condition
Critical	All participating leaf devices report a Major service health state.
Major	One or more VXLAN tunnels to intended VTEPs are not established or are operationally down.
Healthy	All VXLAN tunnels to all intended VTEP peers are established and operationally up.
Spine / Super Spine	Always Healthy — these devices do not participate in tenant service termination.

Warning Threshold Settings

Use the **`efa fabric setting update`** command to configure the link-count thresholds that trigger Warning states. These settings apply fabric-wide.

Triggering a Health Recalculation

Health Check Interval

By default, the health check runs every **30 minutes**. You can change this interval to suit the scale of your deployment and how quickly you want drift to be reconciled. A shorter interval reconciles more frequently; a longer interval reduces background processing on large fabrics.

The interval is a persistent service setting. Once you change it, the new value is stored and remains in effect across service restarts until you change it again.

Interval Format

Specify the interval as a duration string in one of the following forms:

Format	Example	Meaning
Nm	30m	N minutes
Nh	1h	N hours
NhMm	1h30m	N hours and M minutes

For backward compatibility, a plain integer is accepted and interpreted as minutes (for example, 60 is treated as 1h). Values are always stored and displayed in canonical duration form — for example, 90 becomes 1h30m.

Validation Rules

- **Minimum:** 30m (30 minutes).
- **Maximum:** 24h (1440 minutes).
- A value outside the supported range is rejected (CLI error, or HTTP 400 for the REST API).
- An unrecognized format — for example 30s or 24h1m — is rejected.
- Submitting the value that is already stored is allowed and returns success.

How Interval Changes Take Effect

When you set a new interval, the health check is rescheduled immediately and non-disruptively:

- Any health check already in progress runs to completion; only the next scheduled run is affected.
- The schedule restarts from the time of the change and does not carry over timing from the previous schedule.
- The new value is written to persistent storage so it survives a service restart.

Example:
If the interval is 30 minutes and the last run started at 00:30, and you change the interval to 1 hour at 00:45, the next run occurs at 01:45 — one hour after the change.

Note: Setting the interval to the value that is already stored is treated as a no-op success: the setting is accepted, but the schedule is not restarted.

Configure Fabric Health Cron Interval

You can view and update the interval at which the fabric health cron job runs. By default, the fabric health cron interval is set to 30 minutes.

About This Task

The fabric health cron job periodically calculates the health status of the fabric. Follow this procedure to view the current cron interval or update it to a custom value.

Note
Setting a lower cron interval results in more frequent health calculations, which may increase system load. Setting a higher interval reduces load but delays detection of fabric health changes.

Procedure

1. Manage the Interval using the CLI.

   Use the EFA CLI to view and change fabric service settings, including the health check interval.

   a. Display the current settings.

      To list all fabric service settings and their current values, run the following command:

      ```
      efa fabric debug service-settings show
      ```

      Sample output:

      ```
      +----------------------+-------+
      | Property             | Value |
      +----------------------+-------+
      | health_cron_interval | 30m   |
      +----------------------+-------+
      ```

   b. Update the Interval.

      To change the health check interval, run the following command:

      ```
      efa fabric debug service-settings update --health-cron-interval <value>
      ```

      The following flag is supported:

Flag	Required	Format	Range
--health-cron-interval	Yes	30m / 1h / 1h30m	30m – 24h

```
efa fabric debug health cron-interval update --interval 15m
Cron interval updated successfully.
--- Time Elapsed: 1.234567890s ---
```

Note
The updated interval takes effect immediately. You do not need to restart any services.

Examples

```
efa fabric debug service-settings update --health-cron-interval 1h
efa fabric debug service-settings update --health-cron-interval 1h30m
efa fabric debug service-settings update --health-cron-interval 45m
```

Error Messages

The CLI returns a descriptive error in the following cases:

Condition	Message
No flag provided	no settings provided; use -h to see available flags
Value out of range	health_cron_interval must be between 30m and 24h
Unrecognized format	invalid format for health_cron_interval

c. Capture the setting in running-config.

When the interval differs from the built-in default, the update command is included in the EFA running configuration so the deployment state can be captured and replayed. For example:

```
efa fabric debug service-settings update --health-cron-interval 1h30m
```

Note
This line appears in the running configuration only when at least one service setting differs from its default, or when a non-default value has been explicitly stored.

2. Manage the Interval using the REST API.

The service-settings endpoints are exposed under the Fabric Service debug namespace. All requests require a bearer token.

a. Retrieve the current settings.

Attribute	Value
Method	GET
Path	/v1/fabric/debug/service-settings
Authentication	Bearer token required

Response — 200 OK

```
{
  "items": [
    { "key": "health_cron_interval", "value": "30m" }
  ]
}
```

b. Update the settings.

Attribute	Value
Method	PUT
Path	/v1/fabric/debug/service-settings
Content-Type	application/json
Authentication	Bearer token required

- Request body:

```
{
  "items": [
    { "key": "health_cron_interval", "value": "1h30m" }
  ]
}
```

- Response — 200 OK

```
{
  "items": [
    { "key": "health_cron_interval", "value": "1h30m" }
```

```
  ]
}
```

Note
Multiple settings can be submitted in a single request. Updates are applied atomically: if any key in the request is not recognized, the entire batch is rolled back. Submitting a value that is already stored returns 200 OK as a no-op success and does not reschedule the health check.

3. Monitor and verify health checks.

   Each health check records its start and completion in the Fabric Service logs. Use these entries to confirm that checks are running on the expected schedule and to gauge how long a full reconciliation takes.

```
CalculateAllFabricsHealth  Starts health calculation  <start time>
CalculateAllFabricsHealth  Ends health calculation    <end time>
```

   Completion time scales with the number of devices in the fabric. As a reference point, an eight-node fabric completes a full health check in approximately 40 seconds to 1 minute.

Tip
If you shorten the interval on a large fabric, confirm from the start and end log entries that each check completes comfortably within the configured interval, so runs do not overlap.

Fabric Health Failure Scenarios — 3-Clos Topology (Two MCT Pairs)

The following use cases describe expected health outcomes for a 3-clos topology consisting of two spine switches (Spine1, Spine2) and four leaf switches in two MCT pairs (Leaf11/Leaf12

and Leaf21/Leaf22). Link1 and Link2 are MCT connections between the leaf pairs; Link3–Link6 are spine-to-leaf P2P links.

Use Case	Scenario	Health Outcomes	Alerts Generated
1	Spine1 down	Spine1: Minor Spine2: Healthy Leaf11–22: Minor (still connected via Spine2) Fabric: Minor	*FabricPhysicalTopologyDegradedAlert (Spine1, all leaves)* *FabricDeviceUnderlayTopologyDegradedAlert (Spine1, all leaves)* *FabricHealthDegradedAlert*
2	Leaf11 down (local-bias ENABLED)	Leaf11: Major (isolated — device down) Leaf12: Minor (MCT peer degraded, bias handles BUM) Leaf21,22: Minor (overlay partially degraded) Spine1,2: Minor Fabric: Major	*FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricPhysical/Underlay/ OverlayTopologyDegradedAlert + FabricDeviceMctDegradedAlert (Leaf11)* *FabricDeviceOverlayTopologyDegradedAlert + MctDegradedAlert (Leaf12)* *FabricDeviceOverlayTopologyDegradedAlert (Leaf21, Leaf22)* *FabricHealthDegradedAlert*
3	Leaf11 down (local-bias DISABLED)	Leaf11: Major Leaf12: Major (overlay down, no bias fallback) Leaf21,22: Major (overlay down) Spine1,2: Minor Fabric: Major	*FabricPhysical/Underlay/Overlay + MctDegradedAlert (Leaf11)* *FabricDeviceOverlay + MctDegradedAlert (Leaf12, Leaf21, Leaf22)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricHealthDegradedAlert*

Use Case	Scenario	Health Outcomes	Alerts Generated
4	Leaf11 and Leaf12 both down	Leaf11,12: Major (completely isolated) Leaf21,22: Major (no overlay path to MCT pair 1) Spine1,2: Minor Fabric: Major	*FabricPhysical/Underlay/Overlay + MctDegradedAlert (Leaf11, Leaf12)* *FabricDeviceOverlayTopologyDegradedAlert (Leaf21, Leaf22)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricHealthDegradedAlert*
5	Spine1 and Spine2 both down	Spine1,2: Major (entire spine layer isolated) All Leaves: Major (no physical/ underlay path to any spine) Fabric: Critical (all devices Major)	*FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricPhysical/Underlay/ OverlayTopologyDegradedAlert (all leaves)* *FabricHealthDegradedAlert (Critical)*
6	Link1 down (one MCT link; min-mct-link=2)	Leaf11,12: Warning (1 MCT link active < threshold of 2) Leaf21,22: Healthy Spine1,2: Healthy Fabric: Warning	*FabricDeviceMctDegradedAlert (Leaf11, Leaf12)* *FabricHealthDegradedAlert*
7	Link1 & Link2 down (all MCT links; min-mct-link=2)	Leaf11,12: Warning (0 MCT links active < threshold of 2) Leaf21,22: Healthy Spine1,2: Healthy Fabric: Warning	*FabricDeviceMctDegradedAlert (Leaf11, Leaf12)* *FabricHealthDegradedAlert*

Use Case	Scenario	Health Outcomes	Alerts Generated
8	Link3 down (Leaf11↔Spine1 physical + BGP)	Leaf11: Minor Spine1: Minor All others: Healthy Fabric: Minor	*FabricPhysical/ UnderlayTopologyDegradedAlert (Leaf11, Spine1)* *FabricHealthDegradedAlert*
9	Link3 & Link4 down (Leaf11,12↔Spine1)	Leaf11,12: Minor Spine1: Minor Leaf21,22: Healthy Spine2: Healthy Fabric: Minor	*FabricPhysical/ UnderlayTopologyDegradedAlert (Leaf11, Leaf12, Spine1)* *FabricHealthDegradedAlert*
10	Links 3,4,5,6 down (all leaves↔Spine1)	Leaf11,12,21,22: Minor Spine1: Minor Spine2: Healthy Fabric: Minor	*FabricPhysical/ UnderlayTopologyDegradedAlert (all leaves, Spine1)* *FabricHealthDegradedAlert*
11	Links 3,4 (spine) + Links 1,2 (MCT) down	Leaf11,12: Minor (both spine and MCT links down, not isolated) Spine1: Minor Leaf21,22: Healthy Spine2: Healthy Fabric: Minor	*FabricPhysical/Underlay + MctDegradedAlert (Leaf11, Leaf12)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1)* *FabricHealthDegradedAlert*

Use Case	Scenario	Health Outcomes	Alerts Generated
12	Leaf11, Leaf12, Leaf21 down	Leaf11,12,21: Major Leaf22: Major (no overlay peers remaining) Spine1,2: Minor Fabric: Major	*FabricPhysical/Underlay/Overlay + MctDegradedAlert (Leaf11, Leaf12, Leaf21)* *FabricDeviceOverlay + MctDegradedAlert (Leaf22)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricHealthDegradedAlert*
13	All four leaves down	All leaves: Major Spine1,2: Major (no leaf connectivity — spine layer isolated) Fabric: Critical	*FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricPhysical/Underlay/Overlay + MctDegradedAlert (all leaves)* *FabricHealthDegradedAlert (Critical)*

Fabric Health Failure Scenarios — 5-Clos Topology (Two PODs)

This topology adds a super-spine layer (SuperSpine1, SuperSpine2) above two pods, each containing two spines and multiple MCT leaf pairs. Spines in Pod A are Spine1/2; Pod B spines are Spine3/4. Each pod has two MCT leaf pairs.

Note
A 5-Clos fabric becomes Critical when all devices within a single pod reach Major status — not necessarily the entire fabric.

Use Case	Scenario	Health Outcomes	Alerts Generated
1	Spine1 down (5-Clos)	Spine1: Minor Leaf11–22 (Spine1's pod): Minor Super Spine1,2: Minor All other devices: Healthy Fabric: Minor	*FabricPhysical/ UnderlayTopologyDegradedAlert (Leaf11,12,21,22, Spine1, SuperSpine1,2)* *FabricHealthDegradedAlert*
2	Leaf11 down — local-bias ENABLED (5-Clos)	Leaf11: Major Leaf12: Minor (MCT peer, bias handles BUM) All other leaves: Minor (overlay degraded) Spine1,2: Minor (no Leaf11 connectivity) Spine3,4: Healthy Super Spine1,2: Healthy Fabric: Major	*FabricPhysical/Underlay/Overlay + MctDegradedAlert (Leaf11)* *FabricDeviceOverlay + MctDegradedAlert (Leaf12)* *FabricDeviceOverlayTopologyDegradedAlert (Leaf21-42)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricHealthDegradedAlert*
3	Leaf11 down — local-bias DISABLED (5-Clos)	Leaf11: Major Leaf12,21,22,31,32,41,42: Major (no overlay fallback) Spine1,2: Minor Spine3,4: Healthy SuperSpine1,2: Healthy Fabric: Major	*FabricPhysical/Underlay/Overlay + MctDegradedAlert (Leaf11)* *FabricDeviceOverlay + MctDegradedAlert (Leaf12)* *FabricDeviceOverlayTopologyDegradedAlert (Leaf21-42)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricHealthDegradedAlert*

Use Case	Scenario	Health Outcomes	Alerts Generated
4	Leaf11 & Leaf12 down (5-Clos)	Leaf11,12: Major All other leaves: Major (no overlay path) Spine1,2: Minor Spine3,4: Healthy SuperSpine1,2: Healthy Fabric: Major	*FabricPhysical/Underlay/Overlay + MctDegradedAlert (Leaf11, Leaf12)* *FabricDeviceOverlayTopologyDegradedAlert (all remaining leaves)* *FabricPhysical/ UnderlayTopologyDegradedAlert (Spine1, Spine2)* *FabricHealthDegradedAlert*
5	Spine1 & Spine2 down — whole pod spine layer (5-Clos)	Spine1,2: Major Leaf11–22 (pod A): Major Leaf31–42 (pod B): Major (no cross-pod overlay) SuperSpine1,2: Minor (partial connectivity loss) Spine3,4: Healthy Fabric: Critical (entire pod A is Major)	*FabricPhysical/Underlay/ OverlayTopologyDegradedAlert (Leaf11-22, Spine1,2)* *FabricDeviceOverlayTopologyDegradedAlert (Leaf31-42)* *FabricPhysical/ UnderlayTopologyDegradedAlert (SuperSpine1,2)* *FabricHealthDegradedAlert (Critical)*
6	One MCT link Leaf11↔Leaf12 down (min-mct-link=2, 5-Clos)	Leaf11,12: Warning (1 link < threshold 2) All others: Healthy Fabric: Warning	*FabricDeviceMctDegradedAlert (Leaf11, Leaf12)* *FabricHealthDegradedAlert*
7	All MCT links Leaf11↔Leaf12 down (min-mct-link=2, 5-Clos)	Leaf11,12: Minor (0 MCT links, threshold crossed with complete MCT loss) All others: Healthy Fabric: Minor	*FabricDeviceMctDegradedAlert (Leaf11, Leaf12)* *FabricHealthDegradedAlert*
8	Spine1↔Leaf11 link down (5-Clos)	Leaf11: Minor Spine1: Minor All others: Healthy Fabric: Minor	*FabricPhysical/ UnderlayTopologyDegradedAlert (Leaf11, Spine1)* *FabricHealthDegradedAlert*
9	SuperSpine1 down	SuperSpine1: Minor (super-spine layer not fully isolated) SuperSpine2: Healthy Spine1–4: Minor (partial super-spine connectivity lost) All leaves: Healthy Fabric: Minor	*FabricPhysical/ UnderlayTopologyDegradedAlert (SuperSpine1, Spine1-4)* *FabricHealthDegradedAlert*

Use Case	Scenario	Health Outcomes	Alerts Generated
10	SuperSpine1 & SuperSpine2 both down	SuperSpine1,2: Major (complete super-spine isolation) Spine1–4: Minor (super-spine connectivity lost) All leaves: Major (no cross-pod overlay path) Fabric: Major	*FabricPhysical/ UnderlayTopologyDegradedAlert (SuperSpine1,2, all spines)* *FabricDeviceOverlayTopologyDegradedAlert (all leaves)* *FabricHealthDegradedAlert*
11	All spines (Spine1–4) down	Spine1–4: Major SuperSpine1,2: Major (no spine connectivity) All leaves: Major Fabric: Critical (all devices Major)	*FabricPhysical/ UnderlayTopologyDegradedAlert (SuperSpine1,2, Spine1-4)* *FabricDeviceOverlayTopologyDegradedAlert (all leaves)* *FabricHealthDegradedAlert (Critical)*
12	Spine2, Spine3, Spine4 down (Spine1 healthy)	Spine2,3,4: Major Spine1: Healthy SuperSpine1,2: Minor (partial spine connectivity) All leaves: Major (cross-pod overlay down) Fabric: Critical (all devices in at least one pod are Major)	*FabricPhysical/ UnderlayTopologyDegradedAlert (SuperSpine1,2, Spine2,3,4)* *FabricDeviceOverlayTopologyDegradedAlert (all leaves)* *FabricHealthDegradedAlert (Critical)*

Best Practices to Avoid Fabric Health Failure

- Configure appropriate min-mct-link and min-p2p-link thresholds.
- Monitor fabric health regularly through REST APIs or CLI.
- Investigate Major and Critical alerts immediately.
- Maintain redundant spine and super-spine connectivity.
- Validate health status after planned maintenance activities.
- Use manual health calculation to verify topology changes when required.

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.

- Creation of 3-Stage Clos Fabric on page 1045
- Addition of Two Spine Devices 10,20.246.1 and 2 on page 1046
- Addition of Two MCT Pairs 10.20.246.5,6 and 10.20.20.246.3 and 4 on page 1048
- Configuration of Fabric on page 1054

- Shutting Down MCT Port Channel on Leaf Devices on page 1056
- Choosing Not to Shut Down MCT Leaf Nodes on page 1061
- Removing Cluster Config from One MCT Leaf on page 1063
- Triggering DRC to Reconcile Router BGP Config on Switch on page 1090
- Removing Leaf to Spine Links Between Devices on page 1070
- Adding Incorrect Links Between Two Spines on page 1078
- Removing Incorrect Links Between Two Spines on page 1081
- Post Fabric or DRC Configuration on page 1083
- Deleting Router BGP Configuration in a Spine Device on page 1085
- Triggering DRC to Reconcile Cluster Config on page 1066
- Deleting BGP Neighbors from Leaf Device on page 1093
- Configuring Neighbor Again on Switch on page 1099
- Reloading a Leaf Device on page 1101
- Verification of Spine Devices Deletion from Fabric on page 1106
- Removing All Devices from Fabric on page 1108

Creation of 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": {}
}
```


Addition of 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  | Extreme 8720 | 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
-------------------------------------------------------------------------------
```

Addition of 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  | Extreme 8720-32C|
20.4.2slxos20.4.2_220803_1000|      |        |
+----+-------------+-----------+-------+--------------+-------------------------------
+--------+--------+
| 9  | 10.20.246.4 | NH-Leaf2  | 3012  | Extreme 8720-32C|
20.4.2slxos20.4.2_220803_1000|      |        |
+----+-------------+-----------+-------+--------------+-------------------------------
+--------+--------+
| 3  | 10.20.246.5 | NHF-Leaf1 | 3009  | Extreme 8520--48Y| 20.4.1b
|        |        |
+----+-------------+-----------+-------+--------------+-------------------------------
+--------+--------+
| 1  | 10.20.246.6 | NHF-Leaf2 | 3009  | Extreme 8520--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 ---
```

Adding spine and leaf devices turns fabric topology Green, but device health stays Red due to missing device configs (cluster, BGP neighbors). These are device-level issues.

```
(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 |
+------------+-------------+-----------------+-------------+-----------+-------------
+-------------+----------+---------+----------+------------------------+
-------------------------------------------------------------------------------------------------
-------------------------------------------------------------------

=================================================================================================
=================================================================
```

Configuration of 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 ---
```

Shutting Down 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 ---
```

Choosing Not to Shut Down 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 ---
```

Removing 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 ---
```

Triggering 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 ---
```

Removing 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 ---
```

Adding 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 ---
```

Removing 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 Fabric or DRC Configuration

The following is a sample output of 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 ---
```

Deleting 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 ---
```

Triggering 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 ---
```

Deleting 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 ---
```

Configuring Neighbor Again on Switch

The following sample output configures neighbors on 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 ---
```

Reloading a Leaf Device

The following sample output shows the configuration 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 ---
```

Verification of Spine Devices Deletion from 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 ---
```

Removing All Devices from Fabric

The following sample output removes all devices from the 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 ---
```

XCO License Service Management

Starting with the 3.4.0 release, XCO introduces a new licensing service. This service primarily handles offline license checks. Only users with the System Admin role can add licenses to the system. Additionally, users with Fabric Admin or System Admin roles can view the licenses installed in the system.

Using XCO licensing services, you can validate the offline license provided by the users with the following operations:

- Display the license details including features, license start and expiration date, license type, and others.
- Display all the licenses that the users have entered into the system.
- Add or show licenses through CLI and REST API.
- Set alerts to notify you when a license is about to expire. You can configure the threshold for these alerts as needed.
- If necessary, delete licenses from the system.

XCO Licensing Overview

Read the following topics to learn about:

- XCO licensing and its type
- Licensing configuration tasks

XCO-based Licensing Overview

You can purchase XCO licenses as per your workspace (fabric skills) requirements.

A license file can contain multiple licenses, each with its own expiration date and supported features.

The following diagram illustrates the SKUs requirements built on the concept of composable skills:

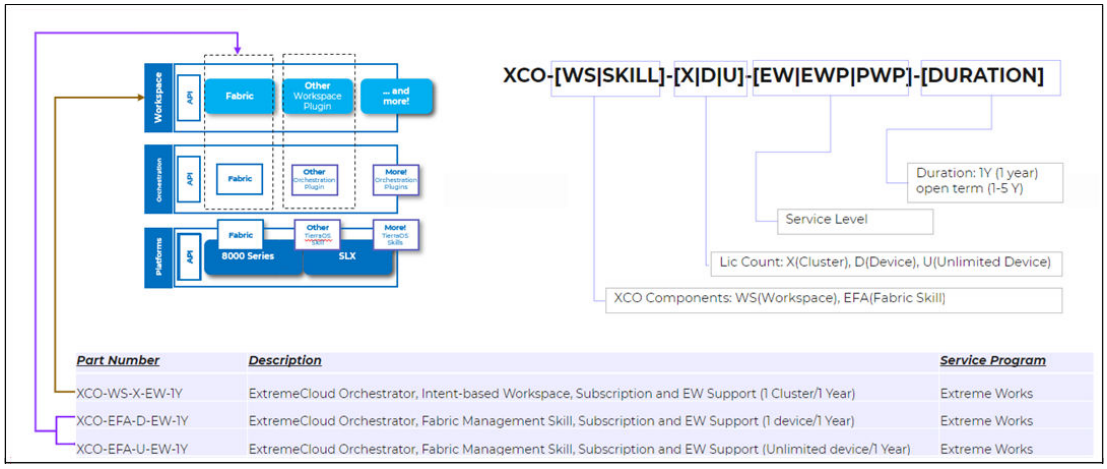

The following is an example XCO license file:

```
21 PRD-XCO-EFA-D Ni LONG NORMAL STANDALONE AGGR 1_KEYS INFINITE_KEYS 12 OCT 2023 0 0 11
NOV 2023 23 59 NiL SLM_CODE CL_ND_LCK NiL *1MTFRGRCYEDRBP20400 NiL NiL NiL 5_MINS NiL 0
0XoTcNq2rzqlqeW04TQ5d8sjR/
ztUtxYVD1aEA12wGoJ+f8bRyV1tGSvelpKuM24U+iADhp24YsLNTIbTnoY9ot2TsHk3Ant12zrglOoLtf8TEY9+Mjl
IjXLCV56of3gTvTV0ssmtuECK8HOhfGZriYvQ1Ckiym7Sp7JV23fqc6KEa6mtRzUke5ljt+CyH5sjxkEVNr6cVTgkc
1AY/Ynch0tYBw/DRO6d/6MqfBCsC/
FPCLSsSG0c5BHI28d9RGNEsKN+Cd7P1eqVKH1wISB7oPupqL8zjFgV6etKnYW6bSZyrL7yrp2C//
jUYuEy4CbPFdXo3sYN3+AAc7NhDX6avfY9EPCN12Yw6cJevPND4VU0CTW2asHhUBVyWUD29y0q1FQ3HgP3qREnXr5f
1Rm54Zodkae+oPX0gA0mdN3xgGrrDRG+sG7Q0aV//l4spUlWNvE/
WwuvfzrIbs9jAOg0R1BOHQTcrTLoFFwZSO5jPEFBFJDQjILgwkFBFJDQjIGAQEHAQEEJThiMWI2ZjFkLTM3NjktNGY
0ZS1iMzg3LTM0ZTMyZWU5NjNiZAAIAQEJQOEpEHKfc/wB+6v/
DsEwOlbbFarKiMWViJSVG4lCC+QJDgLCYiG7kVqNMf/
R9RS9O9fRLjAINmGbiWHnl5+ktdUKgg4wggEKAoIBAQDArKpEU82d/hfQ/AwK/
kzI6A+TndWRiriQ9MGZge2CkQ5AWKi56yrWWBvp032/
qXuWpfIMby9vu9MmU7N1pvsySRUjGsivYRwoQekLUxNiiEMJglBG2/mS9u0/
uli7Lln+pJ6JgTB5O6nKLnJL78qMOb87wkzD4qJfE0pggPfpoLlaWPoKUcth+dnzQ4C2gS2iqL2dwLyBR8WcjleCXH
syJtkS6lBvHOTNgQCFdb6wBIRPqj0IXB8mFKaNCcLRuVCO+IQziyuYYA8p2IFMEGA/
eEriJNS1hYpJLSY2Kmme9+qZY9VkDUtH31rXOsl3pvSyyc89KxEUywy7skZHadCBAgMBAAEEJThiMWI2ZjFkLTM3Nj
ktNGY0ZS1iMzg3LTM0ZTMyZWU5NjNiZAAHAQEMggBsn4aOQtPJdtR6QR8I3pG0bxuw0u/
0W8oro1GgWvjFEAJjEv6POicbCxvZF1clXSbjCOzPW5a+rF2ownaO+tc5ol4mBSwoUZtdcFeoDwSa4p0SWrmAIrPBk
R9P3MOufjLn/
qF8kwPZpgd5nW2TX2crr0/0S6PItYMYtjKEROBpqXHI5gZZ5IbMrEsXn4TloAPdI+z3njBejw7CJ7UxzaAswPEMmdN
TNf+4ywr0JvqCmTNaLt9DIvcaSVh4gphYsmX/
xLFWNDQvnXFp7YkfKuteFxHEgzKFJv5fmUtvur1COPzXP9Omwx6s+U/qw7D6ojXkw6yC/
CW9COhIE+Fvy2vLAAACvw==#RMS 9.7#AID=9074f7f6-dcf5-46b5-8b15-a4370a56807e
```

The following flowchart illustrates licensing process flow in an XCO deployment:

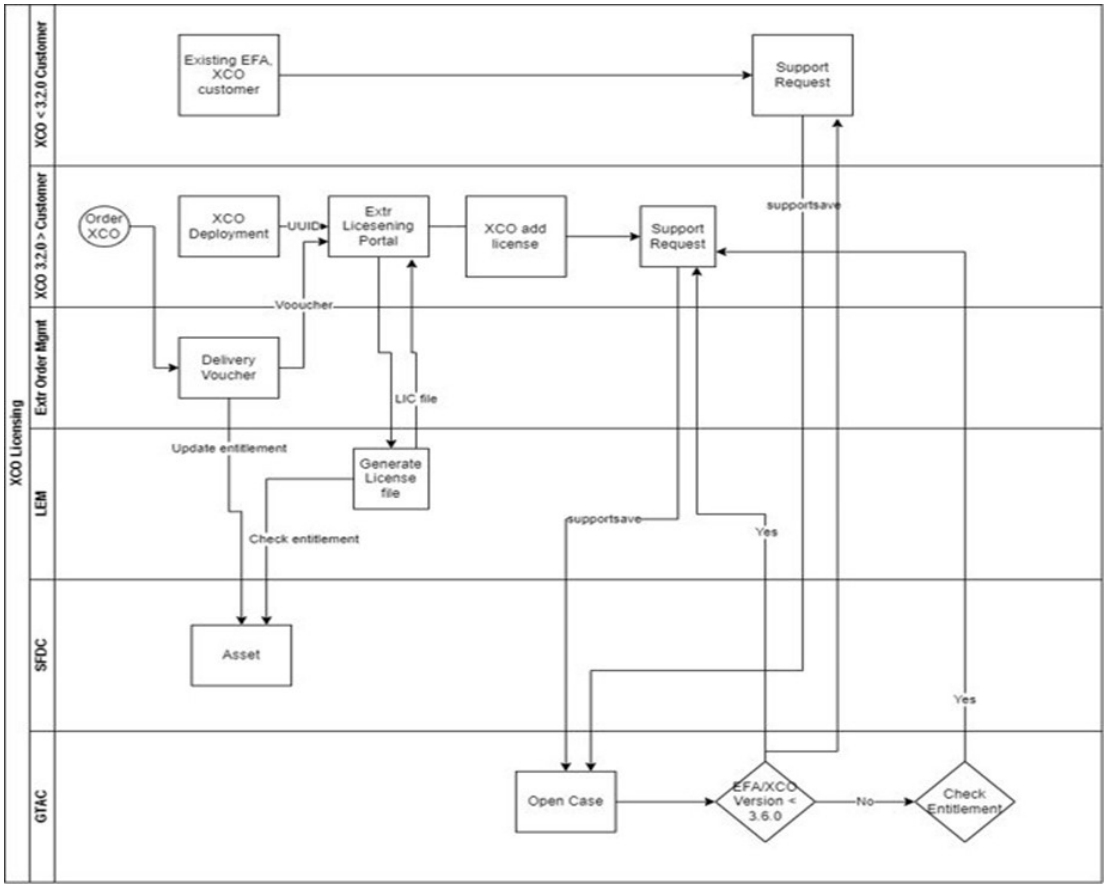

How XCO Licensing Works

A subscription license can be ordered pre-installed in a XCO system when first shipped from the factory, or later ordered and installed by the customer. In either case, additional licenses can be ordered as needed.

When a license is ordered separately (not pre-installed), a fulfillment email message, along with a *Voucher ID (VID)*, is issued to you by Extreme Networks as proof of purchase. The *VID* and *Locking ID (LID)* of the XCO system are used to generate a license key from the Extreme Portal. The license key is contained within a *license file* , which can be uploaded to the switch. You can add the license key to a switch using the **`license add`** command, or the **`license add FTP-URL`** *`ftpPath`* | **`SCP-URL`** *`scpPath`*] command.

XCO License Type

XCO supports subscription licenses.

Instead of making a one-time purchase, you can purchase the products using an annual fee.

XCO is a yearly subscription that has an open term, which can be for a portion of a year. The start and expiry date of the license will reflect this. For instance, if you have XCO-EFA-D-EW-1Y with a 3.5-year open term, the start-expiry duration will be 3.5 years (3 years and 6 months).

XCO License Terminology

The following table lists the terms that are used often through this document:

Term	Definition
License file	The file produced by the Extreme Portal when the license is generated. The file is uploaded to the XCO system and controls access to a *licensed feature* or feature set.
Locking ID (LID)	The identification number that uniquely identifies the XCO system. The LID is used in conjunction with the VID to generate and download a software license from the Extreme Portal. The software license is tied to the LID of the XCO system for which the license was ordered and generated.
Licensed feature	Software feature or set of features that require a valid software license by the XCO system.
Voucher ID (VID)	This unique key, along with the *LID*, is used to generate a software license from the Extreme Portal. The Voucher ID is issued by Extreme Networks when a license is purchased. The Voucher ID is delivered through e-mail which is sent to the customer shortly after the order has completed.

XCO Licensing Configuration Tasks

Learn about the license configuration tasks for generating and obtaining an XCO license, and then installing it on the XCO system.

About This Task
Complete the following configuration tasks for XCO licensing:

Procedure

1. Order the desired license.
   After ordering the license, you receive a fulfillment email containing the Voucher ID.
2. Retrieve the UUID (also known as Locking ID) of the Linux system which XCO is installed on (either on TPVM, or external Ubuntu system), using the **`sudo dmidecode -s system-uuid`** command.
3. Log in to the Extreme Portal located here: https://extremeportal.force.com/ExtrLicenseLanding.
   If you do not have an account yet, log in to https://secure.extremenetworks.com.
4. Upload the license file to the XCO system.
5. Install the license.

Note
Extreme uses the Customer ID and UUID for validating the XCO license entitlement.

Generate a License

You can generate a license from the Extreme Networks Portal.

Before You Begin
Before you can use a software license, you must generate it from the Extreme Networks Portal.

About This Task

Follow this procedure to generate and obtain a software license.

Procedure

1. Log in to the Extreme Networks Portal at https://extremeportal.force.com/ExtrLicenseLanding.
2. If you do not have an Extreme Portal account, select **Create a new Extreme Portal account.**

   The Extreme Portal login window appears:

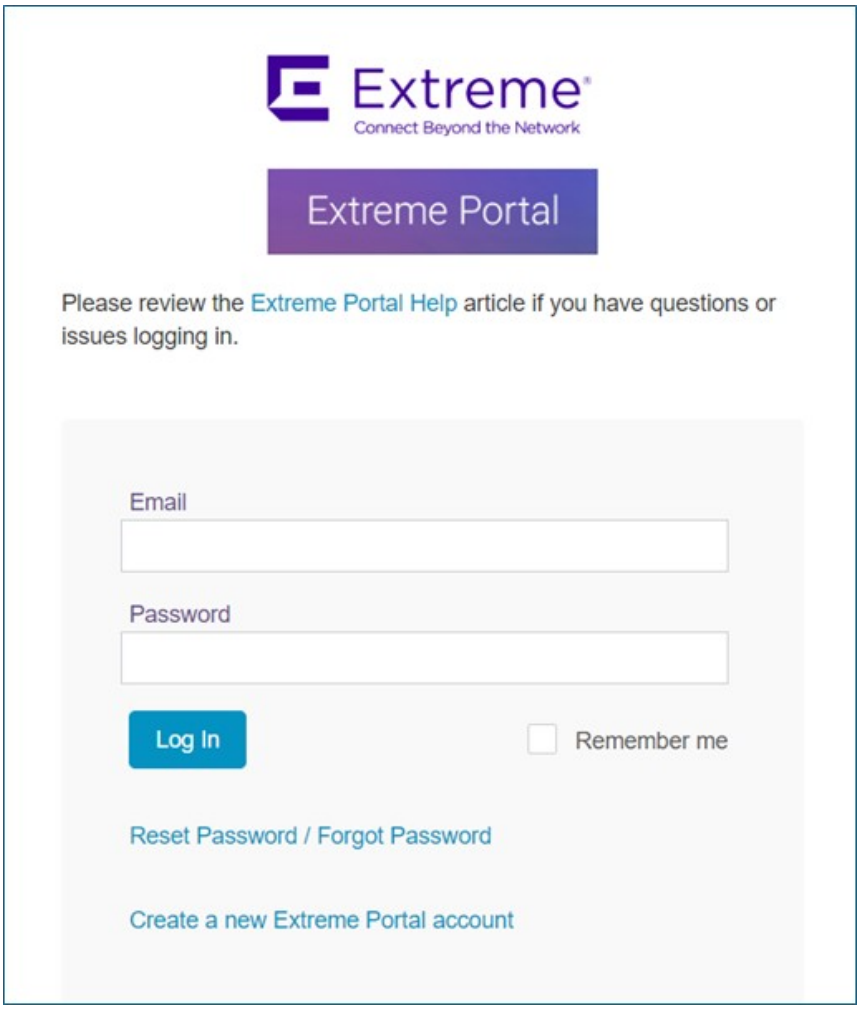

Figure 59: Extreme Portal login window

3. Enter your email and password. Select **Log In**.
   The Extreme Portal home page appears.

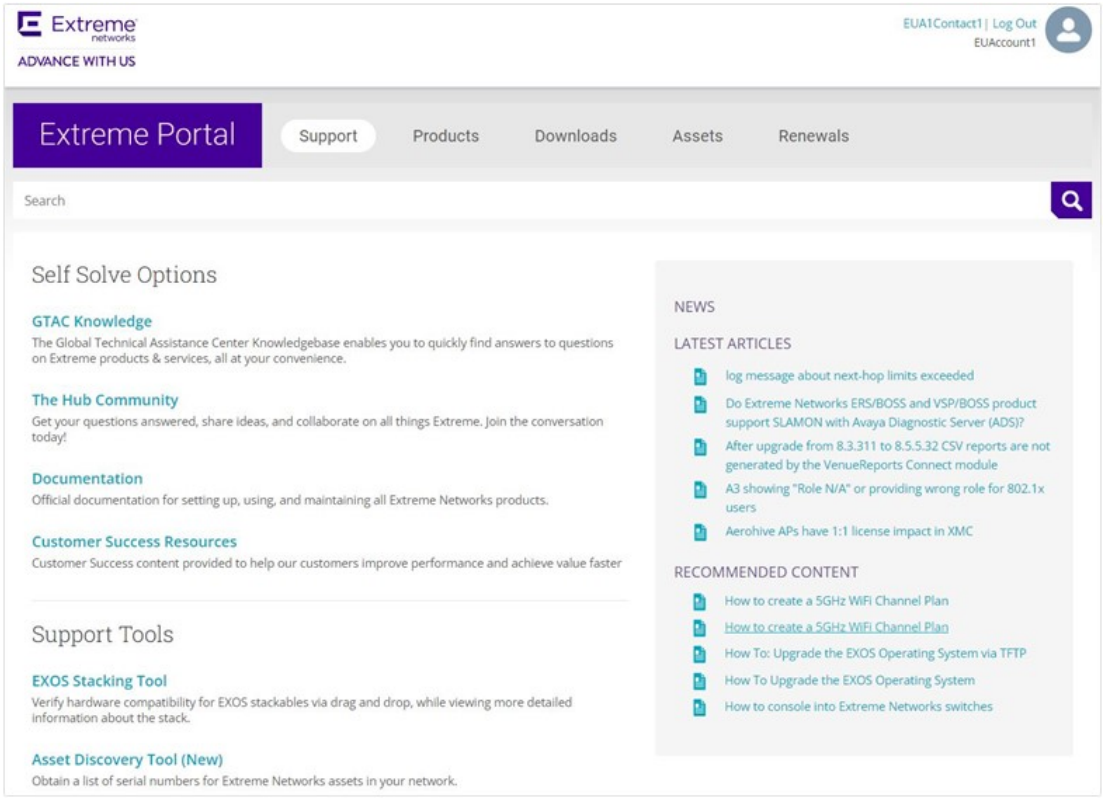

Figure 60: Extreme Portal home page

4. Select the **Assets** tab. Then select **Licenses Home** from the list.
   The **Licenses Home** window appears.

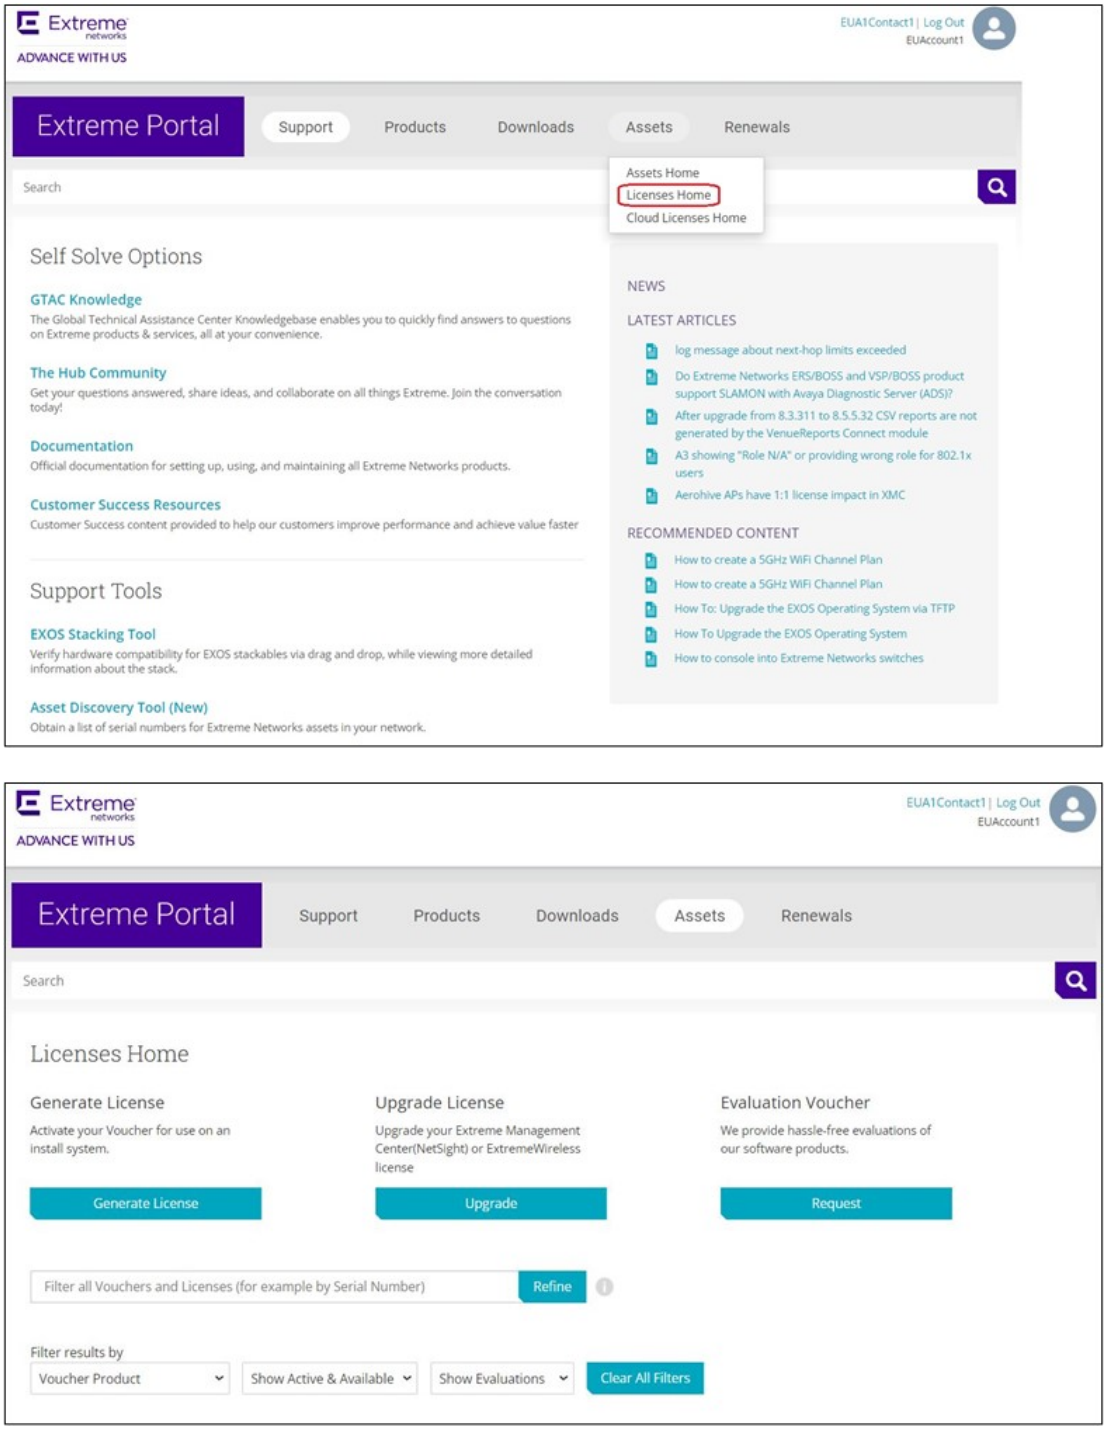

Figure 61: Licenses Home window

5. From the **Licenses Home** page. Select **Generate License**.
   The **Generate License** window appears.

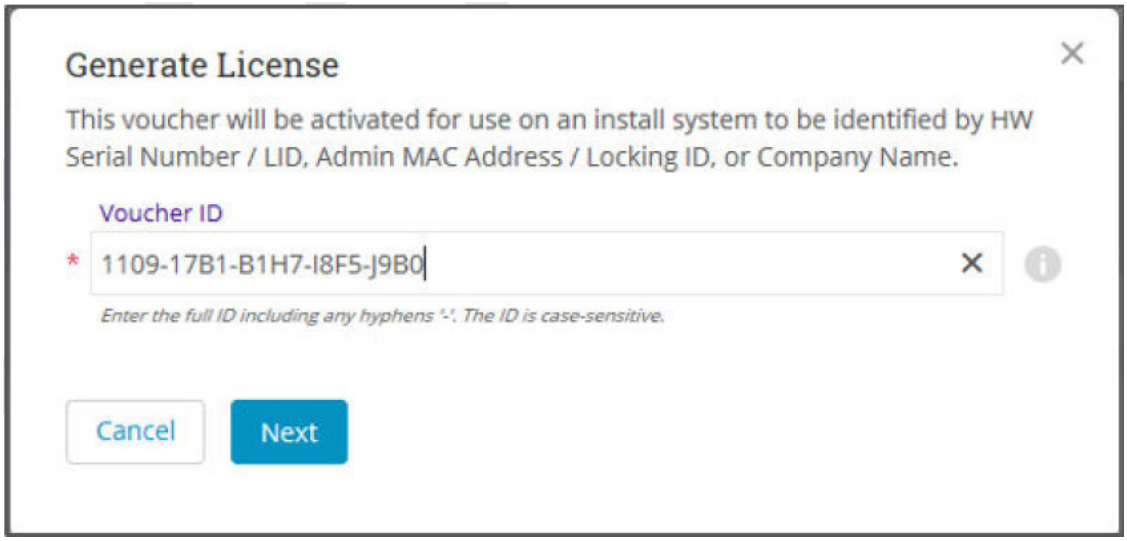

Figure 62: Generate License window

6. Enter the Voucher ID (VID). Then select **Next.**
   The VID contains 20 alphanumeric characters with hyphens.

7. Enter the unique identifier of the HW (or the serial number of the LID). Then select **Submit**.

Note
You must check the box to agree to the Terms and Conditions.

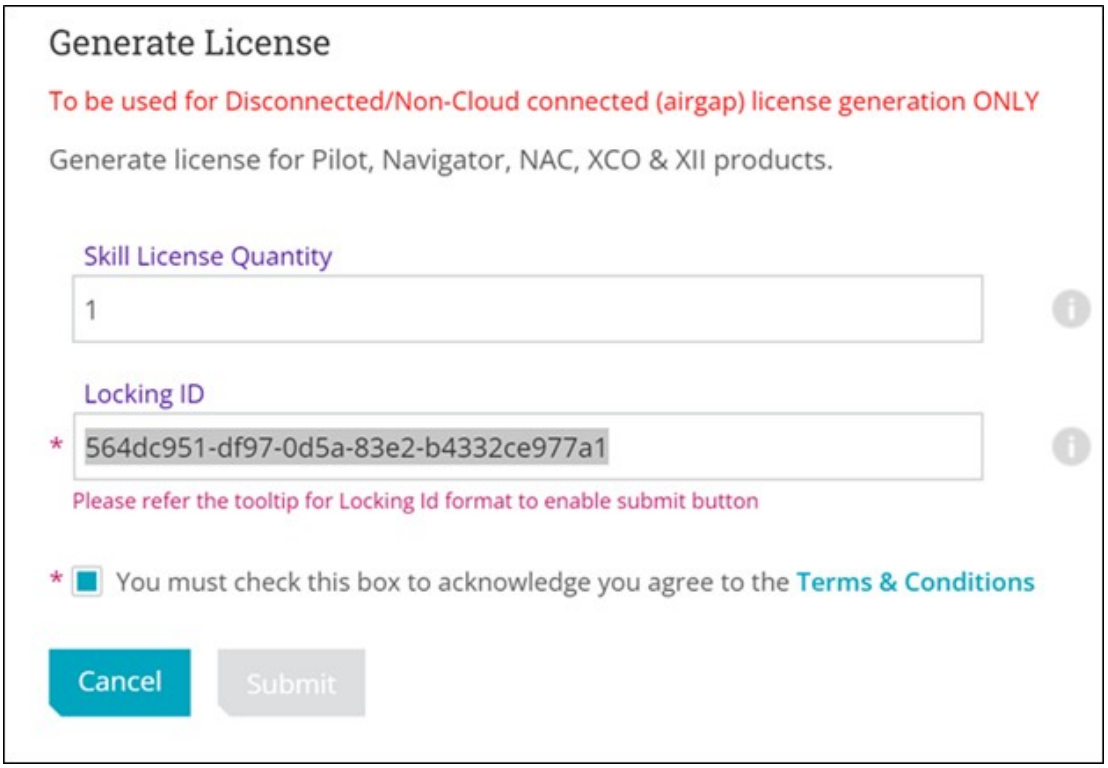

Figure 63: HW serial number window

8. The VID is displayed. The following example shows voucher information for a non-capacity license.

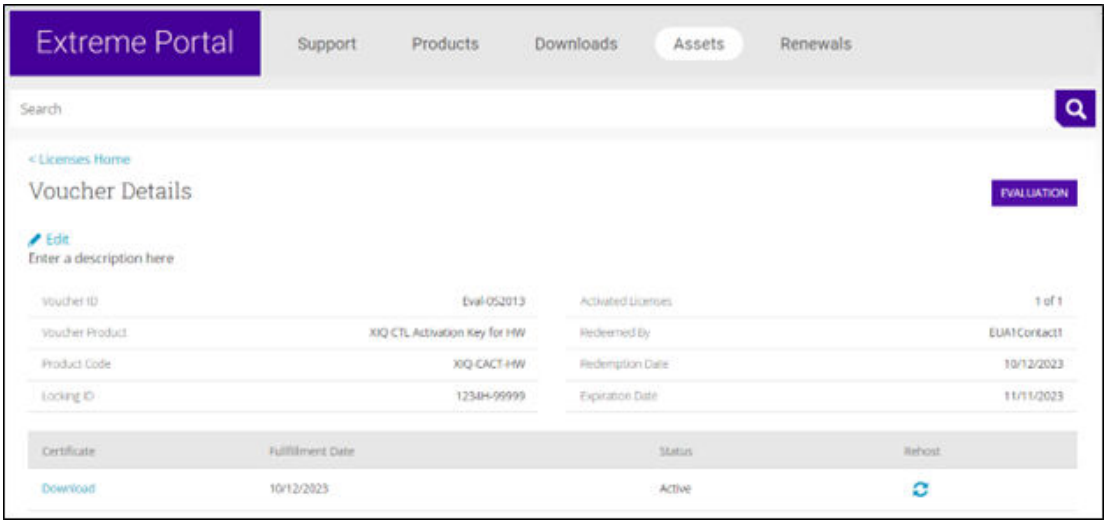

Figure 64: Voucher Details window

9. Select **Download** to download a copy of the license.

   A license can be downloaded after license generation or when a unit is queried. The license is not emailed.

   The following example shows the contents of the XML license, including the license SKU and license keys:

```
21 PRD-XCO-EFA-D Ni LONG NORMAL STANDALONE AGGR 1_KEYS INFINITE_KEYS 12 OCT 2023 0 0
11 NOV 2023 23 59 NiL SLM_CODE CL_ND_LCK NiL *1MTFRGRCYEDRBP20400 NiL NiL NiL 5_MINS
NiL 0 0XoTcNq2rzqlqeW04TQ5d8sjR/
ztUtxYVD1aEA12wGoJ+f8bRyV1tGSvelpKuM24U+iADhp24YsLNTIbTnoY9ot2TsHk3Ant12zrglOoLtf8TEY9+
MjlIjXLCV56of3gTvTV0ssmtuECK8HOhfGZriYvQ1Ckiym7Sp7JV23fqc6KEa6mtRzUke5ljt+CyH5sjxkEVNr6
cVTgkc1AY/Ynch0tYBw/DRO6d/6MqfBCsC/
FPCLSsSG0c5BHI28d9RGNEsKN+Cd7P1eqVKH1wISB7oPupqL8zjFgV6etKnYW6bSZyrL7yrp2C//
jUYuEy4CbPFdXo3sYN3+AAc7NhDX6avfY9EPCN12Yw6cJevPND4VU0CTW2asHhUBVyWUD29y0q1FQ3HgP3qREnX
r5f1Rm54Zodkae+oPX0gA0mdN3xgGrrDRG+sG7Q0aV//l4spUlWNvE/
WwuvfzrIbs9jAOg0R1BOHQTcrTLoFFwZSO5jPEFBFJDQjILgwkFBFJDQjIGAQEHAQEEJThiMWI2ZjFkLTM3Njkt
NGY0ZS1iMzg3LTM0ZTMyZWU5NjNiZAAIAQEJQOEpEHKfc/wB+6v/
DsEwOlbbFarKiMWViJSVG4lCC+QJDgLCYiG7kVqNMf/
R9RS9O9fRLjAINmGbiWHnl5+ktdUKgg4wggEKAoIBAQDArKpEU82d/hfQ/AwK/
kzI6A+TndWRiriQ9MGZge2CkQ5AWKi56yrWWBvp032/
qXuWpfIMby9vu9MmU7N1pvsySRUjGsivYRwoQekLUxNiiEMJglBG2/mS9u0/
uli7Lln+pJ6JgTB5O6nKLnJL78qMOb87wkzD4qJfE0pggPfpoLlaWPoKUcth+dnzQ4C2gS2iqL2dwLyBR8Wcjle
CXHsyJtkS6lBvHOTNgQCFdb6wBIRPqj0IXB8mFKaNCcLRuVCO+IQziyuYYA8p2IFMEGA/
eEriJNS1hYpJLSY2Kmme9+qZY9VkDUtH31rXOsl3pvSyyc89KxEUywy7skZHadCBAgMBAAEEJThiMWI2ZjFkLTM
3NjktNGY0ZS1iMzg3LTM0ZTMyZWU5NjNiZAAHAQEMggBsn4aOQtPJdtR6QR8I3pG0bxuw0u/
0W8oro1GgWvjFEAJjEv6POicbCxvZF1clXSbjCOzPW5a+rF2ownaO+tc5ol4mBSwoUZtdcFeoDwSa4p0SWrmAIr
PBkR9P3MOufjLn/
qF8kwPZpgd5nW2TX2crr0/0S6PItYMYtjKEROBpqXHI5gZZ5IbMrEsXn4TloAPdI+z3njBejw7CJ7UxzaAswPEM
mdNTNf+4ywr0JvqCmTNaLt9DIvcaSVh4gphYsmX/
xLFWNDQvnXFp7YkfKuteFxHEgzKFJv5fmUtvur1COPzXP9Omwx6s+U/qw7D6ojXkw6yC/
CW9COhIE+Fvy2vLAAACvw==#RMS 9.7#AID=9074f7f6-dcf5-46b5-8b15-a4370a56807e
```

Query a License

You can query a license.

About This Task

Follow this procedure to query a license.

Procedure

1. To query a license, select **Assets** > **Licenses Home** option.

2. Enter the SN or LID of the HW asset, or the VID of the SW asset in the **Refine** window box.
   The following example displays how to query a license for a VID.

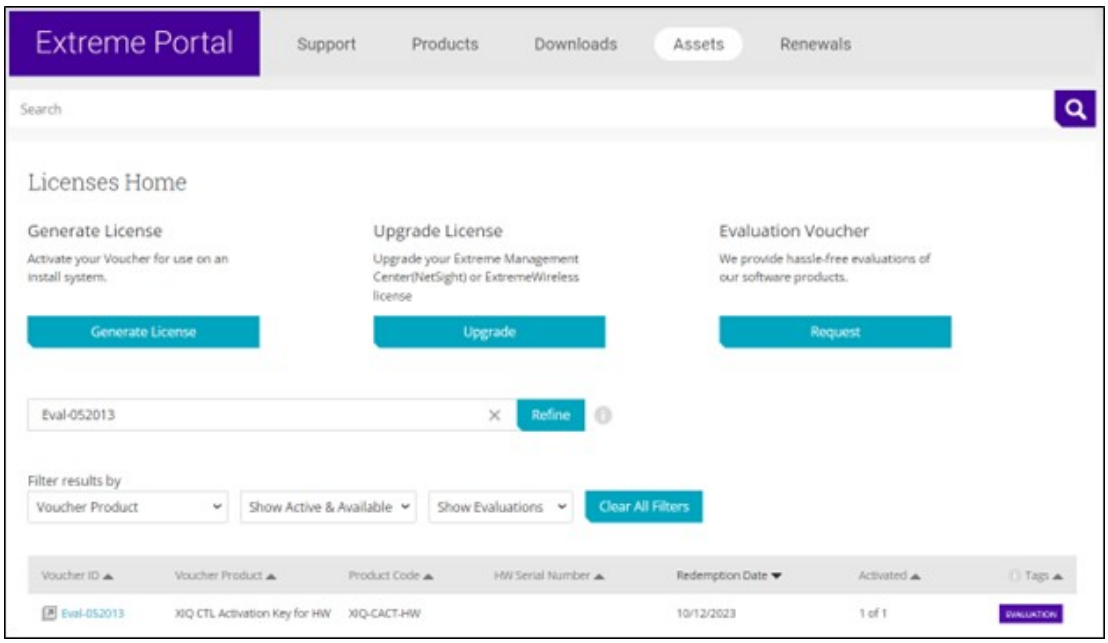

Figure 65: Querying a license window

3. Select the voucher (under the Voucher ID column). The VID is displayed.
   The following example displays voucher information for a non-capacity license.

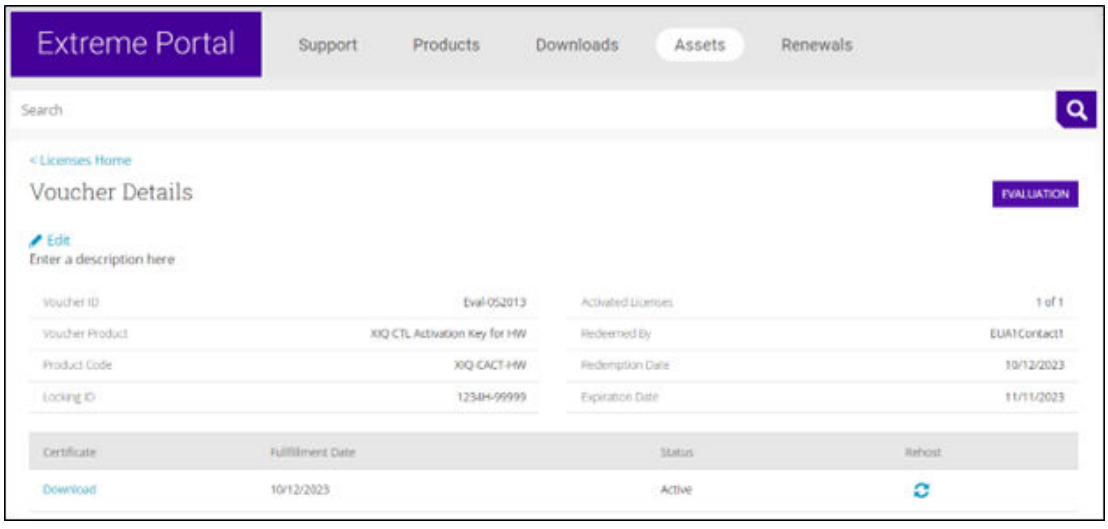

Figure 66: Voucher details window

XCO Licensing Tasks

Learn about all the possible XCO licensing activities.

Install a License

You can install a license on an XCO system.

About This Task

Follow this procedure to install license on an XCO system.

Procedure

1. To install the license on the XCO system, run the following command:

```
efa license add --filepath
```

Note

For information about commands and supported parameters to install a license, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

2. Verify that you added the license by entering the **`show license`** command. The command lists all licensed features currently installed on the XCO system.

Example

```
efa license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic
+-------------------------------------------------------------
+-------------------+-----------------+----------------------------------
+-------------------------------------+
|                          License File
|   License Type    | License Features
|            Start Date             |          Expiration Date           |
+-------------------------------------------------------------
+-------------------+-----------------+----------------------------------
+-------------------------------------+
| lservlcXIQSE-E555555555555555555555555555555551012202312-43AM | Normal Standalone | PRD-
XCO-EFA-D    | at 0:00 hrs (UTC) on Oct 12, 2023 | at 23:59 hrs (UTC) on Nov 11, 2023 |
+-------------------------------------------------------------
+-------------------+-----------------+----------------------------------
+-------------------------------------+
License details
--- Time Elapsed: 976.557741ms ---
```

Configure a License

You can add licenses to the system.

About This Task

Follow this procedure to install and configure license.

For information about commands and supported parameters to configure licenses, see the *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

1. Run the following command to add the licenses:

```
efa license add --filepath <file name with path>
```

2. Run the following command to view the license information:

```
efa license show
```

Example

The following is an example of adding a license:

```
efa license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic
+-------------------------------------------------------------
+------------------+-----------------+----------------------------------
+-----------------------------------+
|                         License File
|   License Type    | License Features
|            Start Date             |          Expiration Date           |
+-------------------------------------------------------------
+------------------+-----------------+----------------------------------
+-----------------------------------+
| lservlcXIQSE-E555555555555555555555555555555551012202312-43AM | Normal Standalone | PRD-
XCO-EFA-D    | at 0:00 hrs (UTC) on Oct 12, 2023 | at 23:59 hrs (UTC) on Nov 11, 2023 |
+-------------------------------------------------------------
+------------------+-----------------+----------------------------------
+-----------------------------------+
License details
--- Time Elapsed: 976.557741ms ---
```

The following is an example of viewing a license information:

```
# efa license show
+-------------------------------------------------------------------------+-------------------
+------------------+------------------+---------------------+
|                             LICENSE FILE                              |   LICENSE TYPE
| LICENSE FEATURES |    START DATE    |   EXPIRATION DATE   |
+-------------------------------------------------------------------------+-------------------
+------------------+------------------+---------------------+
| license1.lic                                                          | Normal Standalone
| PRD-XCO-NAV-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)  |
|                                                                       |
|                  | on Oct 15, 2021   | on Oct 17, 2023     |
| lservlcTEST-0000-0001PRD-5000-PRMR.lic                                | Normal Standalone
| PRD-5000-PRMR    | at 0:00 hrs (UTC) | License has no      |
|                                                                       |
|                  | on Jun 2, 2022    | expiration.         |
| lservlcXCO-4464725328CD4130B8D425187D410000-100xNav-17-10-2023.lic  | Normal Standalone
| PRD-XCO-NAV-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)  |
|                                                                       |
|                  | on Oct 15, 2021   | on Oct 17, 2023     |
| lservlcXIQSE-1112223334445556667890ABCD87878791120231-27PM.lic      | Normal Standalone
| PRD-XCO-PIL-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)  |
|                                                                       |
|                  | on Jan 1, 2023    | on Dec 31, 2023     |
+-------------------------------------------------------------------------+-------------------
+------------------+------------------+---------------------+
License details
--- Time Elapsed: 155.550115ms ---
```

Display a License

You can display installed licenses using the **`show license`** command.

About This Task
Follow this procedure to display a license.

Procedure

Run the following command to view the license information:

```
# efa license show
+-------------------------------------------------------------------------+-------------------
+------------------+------------------+---------------------+
|                            LICENSE FILE                                 |   LICENSE TYPE
| LICENSE FEATURES |    START DATE    |   EXPIRATION DATE   |
+-------------------------------------------------------------------------+-------------------
+------------------+------------------+---------------------+
| license1.lic                                                            | Normal Standalone
| PRD-XCO-NAV-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)  |
|                                                                         |
|                  | on Oct 15, 2021   | on Oct 17, 2023     |
| lservlcTEST-0000-0001PRD-5000-PRMR.lic                                  | Normal Standalone
| PRD-5000-PRMR    | at 0:00 hrs (UTC) | License has no      |
|                                                                         |
|                  | on Jun 2, 2022    | expiration.         |
| lservlcXCO-4464725328CD4130B8D425187D410000-100xNav-17-10-2023.lic  | Normal Standalone
| PRD-XCO-NAV-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)  |
|                                                                         |
|                  | on Oct 15, 2021   | on Oct 17, 2023     |
| lservlcXIQSE-11122233344445556667890ABCD87878791120231-27PM.lic      | Normal Standalone
| PRD-XCO-PIL-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)  |
|                                                                         |
|                  | on Jan 1, 2023    | on Dec 31, 2023     |
+-------------------------------------------------------------------------+-------------------
+------------------+------------------+---------------------+
License details
--- Time Elapsed: 155.550115ms ---
```

When there are no licenses installed on the XCO systems, the **`efa show license`** command output displays the following:

```
efa show license
No entries available
```

Delete a License

You can delete an installed license using the **`efa license delete`** command.

About This Task
Follow this procedure to delete a license from the XCO system.

Procedure

Run the following command to delete the license:

```
# efa license add --
filepath=lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic
 +-------------------------------------------------------------------------
+------------------+-----------------+-----------------------------------
+-----------------------------------+
 | License File | License Type | License Features | Start Date | Expiration Date | AID |
```

```
 +-------------------------------------------------------------------------------
+------------------+------------------+-----------------------------------
+-------------------------------------+
 | lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1
| Normal Standalone | PRD-XCO-EFA-D | at 0:00 hrs (UTC) on
Oct 12, 2023 | at 23:59 hrs (UTC) on Oct 11, 2024 | 5272f34d-e1fd-488c-bf6c-7e2b2d831db1
 +-------------------------------------------------------------------------------
+------------------+------------------+-----------------------------------
+-------------------------------------+
License details
 --- Time Elapsed: 108.919628ms ---

 # efa license delete --
filename=lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic
 License file with AID '5272f34d-e1fd-488c-bf6c-7e2b2d831db'
and filename 'lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-
XCO-EFA-D1.lic' deleted successfully.
 --- Time Elapsed: 53.849675ms ---
```

Licensed Features and Part Numbers

ExtremeCloud Orchestrator uses licenses stored locally in a license file. This ensures XCO does not require an internet connection to verify licenses are available as you add XCO systems.

The following table lists the SKUs available for the customers to purchase for the XCO systems:

XCO-EFA-D-EW-1Y

Description / Product Name	Service Type	Product Family	Product Note	Service Class
ExtremeCloud Orchestrator, Fabric Management Skill, Subscription and EW Support (1 device/1 Year)	XCO On-Prem, Fabric Management Skill, EW Support	XCO		Subscription

XCO-WS-X-EW-1Y

Description / Product Name	Service Type	Product Family	Product Note	Service Class
ExtremeCloud Orchestrator, Intent-based Workspace, Subscription and EW Support (1 Cluster/1 Year)	XCO On-Prem, EW Support	XCO		Subscription

XCO-EFA-U-EW-1Y

Description / Product Name	Service Type	Product Family	Product Note	Service Class
ExtremeCloud Orchestrator, Fabric Management Skill, Subscription and EW Support (Unlimited device/1 Year)	XCO On-Prem, Fabric Management Skill, EW Support	XCO		Subscription

Licensing for XCO Systems in Air Gap Mode

ExtremeCloud Orchestrator uses licenses stored locally in a license file.

For licensing Air Gap mode, an active internet connection is not required for license validation by XCO.

Licensing Supportsave Details

All the decoded licenses are saved in `licenses.txt` and included in the supportsave.

The `licenses.txt` contains output from the **`show licenses`** command.

Note

License service logs are included in supportsave for GTAC to verify the entitlements against Extreme Licensing system. Since no database is utilized by the license service, there is no db dump added to the supportsave file for the licensing functionality.

License Backup and Restore

All licenses are stored in the `<efadatadir>/licenses` file.

The restore function does not change since the files will not be backed up during the backup. In this way, it ensures that a backup from one server can be restored to another, but licenses are not changed.

Note

When you run any upgrade or node-replacement procedure, the license directory remains same and the licenses are retained.

Troubleshooting Licensing Issues

Some features require licenses in order to work properly. Before you call to XCO support, ensure that you have the correct licenses installed on the XCO systems by using the **`show license`** command.

License is Not Properly Installed

If the license is not present or has not been installed correctly, warnings will be issued.

About This Task
Follow this procedure if you suspect a license is not properly installed:

Procedure

1. Run the **`show license`** command to display the currently installed licenses.

```
# efa license show
+-------------------------------------------------------------------------
+-------------------+-----------------+-------------------+--------------------+
|                            LICENSE FILE
|   LICENSE TYPE    | LICENSE FEATURES |    START DATE     |  EXPIRATION DATE  |
+-------------------------------------------------------------------------
+-------------------+-----------------+-------------------+--------------------+
| license1.lic                                                                |
Normal Standalone | PRD-XCO-NAV-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)|
|
|                   |                 | on Oct 15, 2021   | on Oct 17, 2023   |
| lservlcTEST-0000-0001PRD-5000-PRMR.lic                                      |
Normal Standalone | PRD-5000-PRMR    | at 0:00 hrs (UTC) | License has no    |
|
|                   |                 | on Jun 2, 2022    | expiration.       |
| lservlcXCO-4464725328CD4130B8D425187D410000-100xNav-17-10-2023.lic   |
Normal Standalone | PRD-XCO-NAV-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)|
|
|                   |                 | on Oct 15, 2021   | on Oct 17, 2023   |
| lservlcXIQSE-111222333444555666789OABCD87878791120231-27PM.lic       |
Normal Standalone | PRD-XCO-PIL-S-C  | at 0:00 hrs (UTC) | at 23:59 hrs (UTC)|
|
|                   |                 | on Jan 1, 2023    | on Dec 31, 2023   |
+-------------------------------------------------------------------------
+-------------------+-----------------+-------------------+--------------------+
License details
--- Time Elapsed: 155.550115ms ---
```

2. If the license appears in the **`show license`** command output, but the feature does not work for the XCO systems, then you need to reboot the XCO.
3. If the license does not appear in the **`show license`** command output, then it was not installed. Run the **`efa license add --filepath`** command to install the license.

```
efa license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/
lservlcXIQSE-E555555555555555555555555555555551012202312-43AM.lic
+-------------------------------------------------------------
+-------------------+-----------------+---------------------------------
+-----------------------------------+
|                        License File
|   License Type    | License Features
|            Start Date             |          Expiration Date           |
+-------------------------------------------------------------
+-------------------+-----------------+---------------------------------
+-----------------------------------+
| lservlcXIQSE-E555555555555555555555555555555551012202312-43AM |
Normal Standalone | PRD-XCO-EFA-D    |
at 0:00 hrs (UTC) on Oct 12, 2023 | at 23:59 hrs (UTC) on Nov 11, 2023 |
+-------------------------------------------------------------
+-------------------+-----------------+---------------------------------
+-----------------------------------+
License details
--- Time Elapsed: 976.557741ms ---
```

License Error Handling

Make use of the details in this section to learn about all potential error situations when adding licenses.

License File Already Exists

An attempt to add an already-existing license fails with the following error:

```
license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic
CLI Failed.
Error Message:     license file /licenses/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic already exists
Error:             409 Conflict
```

License Expired

An attempt to add an expired license fails with the following error:

```
license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic
CLI Failed.
Error Message:     license already expired on <date>
Error:             403 Forbidden
```

License Parsing Error

An attempt to add a damaged or tampered license fails with the following error:

```
license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic
CLI Failed.
Error Message:     license file cannot be parsed, <error>
Error:             409 Conflict
```

Not an XCO License

An attempt to add a non-XCO license fails with the following error:

```
license add --filepath /Users/SSAVLA/git/GoDCApp/XCO/licensing/scripts/lservlcXIQSE-
E555555555555555555555555555555551012202312-43AM.lic
CLI Failed.
Error Message:     invalid - not an XCO license
Error:             400 Bad Request
```

License Expiry Alert

Alerts are generated before the expiration date of a license. For instance, alerts are generated 90, 60, or 30 days before the expiration date.

The licensing service runs a daily cron job to check for the license expiry. If the expiry occurs within the 90, 60, or 30 days window, a message appears on the message bus. Fault management service receives the message, and then create the alert.

The following alerts are raised based on certain conditions:

License Alerts	Conditions
LicenseExpiryThresholdAlert	Raised when the license expiry date has reached a certain threshold (90 days, 60 days, or 30 days).
LicenseExpiredAlert	Raised when the license has reached its expiry date (day of) and generated every day until license renewal.
LicenseExpiryClearAlert	Raised when the license is renewed and is more than 90 days from expiry.

The following are the severity of alerts based on certain conditions:

- Warning: License is 90 days away from expiry.
- Minor: License is 60 days away from expiry.
- Major: License is 30 or less days from expiry.
- Critical: License is expiring today.
- Info: License is renewed and is more than 90 days away from expiry or license is deleted.

For more details on licensing alerts, see License Alerts on page 989.

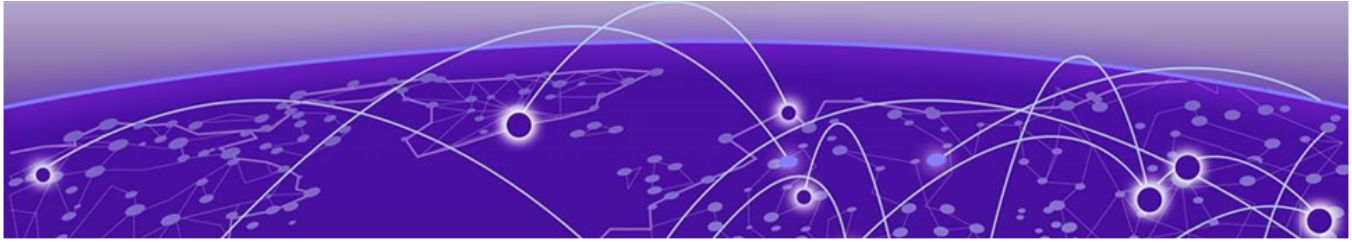

OS ONE Integration with XCO

Learn about Extreme OS ONE configuration and integration with XCO.

Extreme OS ONE

Extreme OS ONE consists of OS software and microservices that provide various services and functionality. This is a new operating system for IP fabrics and data centers.

Extreme OS ONE architecture comprises various components composed as microservices, such as LLDP microservice and BFD microservice. Extreme OS ONE microservices are grouped into two major categories:

- Base OS - comprises the application that is essential for deployment
- Application OS - comprises services such as Switching and Routing, and SDWAN

The Application OS and Extreme OS ONE Switching and Routing provide switching and routing capabilities.

Extreme OS ONE is a cloud-native network operating system (NOS) based on a micro-services architecture. Key characteristics include:

- Modular and composable design for a simple software life-cycle management.
- API-first approach for management programmability.
- Data plane abstraction, supporting integration with multiple ASIC vendors and accelerating the introduction of new hardware platforms.
- Security-first principles that enhance responsiveness to vulnerabilities and minimize the attack surface.

Extreme OS ONE is a high-performance network operating system designed for data centers, service provider, and enterprise networking environments. Extreme OS ONE powers Extreme 8000 series switches and routers.

The build script builds images for all modules and forms an ONIE-compatible bin image during installation for net-install and firmware update. You can download the Extreme OS ONE image from a remote server using any of the following methods:

- FTP
- TFTP

Note

- The XCO Policy Service supports only QoS maps for Extreme OS ONE.
- For Drop Precedence configuration on Extreme OS ONE, see Configure QoS Maps with Drop Precedence on Extreme OS ONE on page 702.
- XCO 4.0.2 does not support Brownfield Auto VNI Offset configuration under overlay networks. If you configure an offset under overlay networks, it will mess up tenant network VNI allocations, causing XCO and OS ONE device VNI values to mismatch devices.

Supported Hardware

Extreme OS ONE Switching and Routing supports the following device:

- 8730-32D (running Extreme OS ONE)

Operations to Integrate OS ONE with XCO

The following are the high-level operations required to integrate OS ONE with XCO:

- Register Inventory Device
- Configure Inventory Device Interface Network Essentials
- Configure Device Level Network Essentials
- Configure Device Level Settings
- Configure IP Fabric BGP-EVPN Single Rack Non-Clos LCM
- Configure Tenant Services

XCO OS ONE Integration and Platform Expansion

XCO 4.0.2 and later releases support for XCO OS ONE devices and adds compatibility with the Extreme 8730, 8720, and 8820 platforms.

Note

- The features related to XCO OS ONE and Extreme 8730 support in XCO 4.1.0 are released as Control Released Features.
- For comprehensive feature descriptions, CLI command references, administrative procedures, and API documentation, contact Extreme Networks.

Extreme OS ONE Overview

Extreme OS ONE is a cloud-native, microservices-based network operating system designed for IP fabrics and data center environments. Key features include:

- **Modular Architecture**: Separates Base OS and Application OS for flexible deployment
- **API-First Design**: Supports gNMI/gNOI with OpenConfig YANG models and CLI.
- **Security-First Principles**: Includes AAA integration, X.509 certificate management, and secure access protocols
- **Lifecycle Management**: Enables firmware install, rollback, commit, diagnostics, and ZTP
- **Monitoring & Logging**: Offers extensive telemetry, syslog support, and threshold-based alerts
- **Integrated Appliance Hosting**: Supports third-party VM deployment via KVM

Supported Platforms

Extreme 8730 (running Extreme OS ONE)

Extreme 8730 Platform

The Extreme 8730-32D is a high-performance switching and routing platform powered by Extreme OS ONE. Its capabilities include:

- **Port Operations**: 400G support with breakout, RS-FEC, and link diagnostics
- **Layer 2 Features**: MLAG, LAG, bridge domains with VLAN tagging and MAC learning
- **Layer 3 Routing**: Full IPv4/IPv6 stack, BGP (iBGP/eBGP), ECMP, and VXLAN/EVPN overlays
- **Resiliency**: BFD, resource threshold monitoring, and thermal/fan policy management

Key XCO Integration Capabilities

XCO offers the following key integration capabilities:

- Device registration and inventory management.
- Secure settings enforcement via gNMI.
- Interface and system configuration.
- Fabric Integration for IP Fabric deployments.
- Drift and reconcile operations

OS ONE Feature Support in XCO

Use this topic to learn about the features introduced in XCO as part of the OS ONE support:

Device Management

- Interface Inventory Enhancements
- Device and Network Settings Inventory
- Switch Lifecycle Management (LCM)
- Switch Firmware Upgrade Support
- Support Save Operations

- Switch Config Backup/Replay
- Device Password Expiry and Clear

Non-Clos Fabric and Network Management

- Non-Clos Fabric Support
- Auto VNI Mapping for BD-Based Tenants
- Auto VNI Mapping for VLAN-Based Tenants
- L2 and L3 EPG Support (BD and VLAN Based)
- VRF Support
- BGP Peer and Peer-Group Support

Health and Monitoring

- RASlog Notifications
- Threshold Monitoring
- SNMP Support and Notifications
- Configuration Tracking
- Switch Health Management (SHM) and DRC
- Alarms, Events, and Fault Health (Phase 1)
- Maintenance Mode Support

Tenant Operations

- Port-Channel (PO) Support
- Auto VNI Mapping for L2/L3 EPGs
- Tenant-Level VRF and BGP Support
- Secure Settings and KV Store Integration
- Compact Fabric QoS Support (Phase 1)
- RMA Support

TPVM Support

- TPVM Support (Single Node and Multinode)
- TPVM Upgrade Support

Other Features

- NTP Support
- Certificate Management
- SNMP Notification Integration

Register Inventory Device

You can register a device in XCO inventory.

Before You Begin

Ensure that you have already deployed XCO.

For detailed procedure on XCO deployment, see *ExtremeCloud Orchestrator Deployment Guide, 4.1.0*.

About This Task

Follow this procedure to register a device in XCO inventory.

Note

For more information on syntax and command examples, see *ExtremeCloud Orchestrator Command Reference, 4.1.0*.

Procedure

1. To register a device in XCO, run the following command:

```
efa inventory device register --ip <list of device-ips> --username admin --password
password
```

Example:

```
efa inventory device register --ip 10.64.224.12-13 --username=admin --password
Rocks@123
+----+--------------+-----------+----------+--------------+----------------------
+---------+---------+
| ID |  IP Address  | Host Name |  Model   | Chassis Name
|       Firmware       | Status  | Reason |
+----+--------------+-----------+----------+--------------+----------------------
+---------+---------+
| 21 | 10.64.224.12 | 40c       | 8820-40C | 8820-40C     |
OSOneSR-22.2.2.0-070 | Success |         |
+----+--------------+-----------+----------+--------------+----------------------
+---------+---------+
| 23 | 10.64.224.13 | 40c       | 8820-40C | 8820-40C
| OSOneSR-22.2.2.0-070 | Success |         |
+----+--------------+-----------+----------+--------------+----------------------
+---------+---------+
Device Details
--- Time Elapsed: 1m40.719295828s ---
(efa:extreme)extreme@TPVM12:~$
```

2. To update a device or to get the current state of a device, run the following command:

```
efa inventory device update --ip <deviceIp>
```

Example:

```
efa inventory device update --ip 10.64.216.44
+----+--------------+-----------+-------+----------------+----------+-----+------
+--------+---------+---------+
| ID |  IP Address  | Host Name | Model |  Chassis Name
| Firmware | ASN | Role | Fabric | Status  | Reason |
+----+--------------+-----------+-------+----------------+----------+-----+------
+--------+---------+---------+
| 15 | 10.64.216.44 | oneos44     | 3200  | Extreme 8520-48Y-8C
| 22.1.3.0 |     |      |        | Success |        |
+----+--------------+-----------+-------+----------------+----------+-----+------
+--------+---------+---------+
Device Details
```

3. To delete a device from XCO inventory, run the following command:

   The device retains its GRPC configuration, JWT token-validator settings, and certificates even after it is deleted from XCO.

   ```
   efa inventory device delete --ip <deviceIp>
   ```

   Example:

   ```
   efa inventory device delete --ip 10.64.216.44
   +----+--------------+-----------+-------+----------------+----------+--------+---------
   +--------+
   | ID |  IP Address  | Host Name | Model |  Chassis Name  | Firmware | Fabric | Status
   | Reason |
   +----+--------------+-----------+-------+----------------+----------+--------+---------
   +--------+
   | 15 | 10.64.216.44 | oneos44     | 3200  | Extreme 8520-48Y-8C | 22.1.3.0 |        |
   Success |        |
   +----+--------------+-----------+-------+----------------+----------+--------+---------
   +--------+
   Device Details
   --- Time Elapsed: 10.621680338s ---
   ```

4. To view the running-config of all current XCO configurations for core services, run the following command:

   ```
   efa show-running-config
   ```

   Example:

   ```
   (efa:extreme)extreme@node-1:~$ efa show-running-config
   efa inventory device secure settings enable
   efa inventory device secure settings update --min-tls-version 1.3 --mac-algorithm hmac-
   sha1,hmac-sha1-96,hmac-sha2-256,hmac-sha2-256-etm-openssh,hmac-sha2-512,hmac-sha2-512-
   etm-openssh --cipher aes128-ctr,aes128-cbc,aes128-gcm-openssh,aes256-ctr,aes192-
   ctr,aes256-cbc,aes256-gcm-openssh,chacha20-poly1305-openssh --key-exchange-algorithm
   curve25519-sha256,curve25519-sha256-libssh,diffie-hellman-group-exchange-sha1,diffie-
   hellman-group-exchange-sha256,diffie-hellman-group14-sha1,diffie-hellman-group14-
   sha256,diffie-hellman-group16-sha512,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-
   nistp521 --telnet Enable --max-password-age 90 --force-default-password-change Enable
   efa inventory device register --ip "10.64.224.17" --username admin --password
   '<encrypted>'
   efa inventory device setting update --ip "10.64.224.17" --maint-mode-enable-on-reboot
   Yes --maint-mode-enable No --mct-bring-up-delay 300 --prefix-independent-convergence
   Yes --crypto-cert-expiry-info 70 --crypto-cert-expiry-minor 50 --crypto-cert-expiry-
   major 25 --crypto-cert-expiry-critical 5 --password-expiry-info 70 --password-expiry-
   minor 50 --password-expiry-major 25 --password-expiry-critical 5 --health-check-enable
   Yes --health-check-interval 6m --health-check-heartbeat-miss-threshold 3 --config-
   backup-periodic-enable Yes --config-backup-interval 2h --number-of-config-backups 5
   efa inventory device interface set-breakout --ip "10.64.224.17" --if-type eth --if-
   name 0/8 --mode 4x100g
   ...
   efa inventory device interface set-description --ip "10.64.224.17" --if-type eth --if-
   name 0/2:1 --description XCO OSONE Interface
   ...
   efa tenant create --name "Tnt-5" --type private --vlan-range 3151-3162,3171,3176 --vrf-
   count 2 --port 10.64.224.16[0/20:2],10.64.224.17[0/20:2]
   efa tenant po create --tenant "Tnt-5" --name "po501" --description " po501" --min-link-
   count 1 --mtu 9000 --negotiation active --number 140 --lacp-timeout short --port
   10.64.224.16[0/20:2],10.64.224.17[0/20:2]
   efa tenant vrf create --tenant "Tnt-5" --name "Tnt-5_VR1" --routing-type "distributed"
   --layer3-extension-enable false --rt-type import --rt 102:102 --rt-type export --rt
   102:102 --ipv4-static-route-next-hop 10.64.224.17,2.4.1.0/24,2.2.1.3 --ipv4-static-
   route-next-hop 10.64.224.17,2.4.1.0/24,2.2.6.4 --ipv4-static-route-bfd
   10.64.224.17,2.2.1.3,,300,300,3 --ipv4-static-route-bfd
   10.64.224.17,2.2.6.4,,300,300,3 --ipv4-static-route-next-hop
   10.64.224.16,2.4.1.0/24,2.3.1.3 --ipv4-static-route-next-hop
   10.64.224.16,2.4.1.0/24,2.3.6.4 --ipv4-static-route-bfd
   ```

```
10.64.224.16,2.3.1.3,,300,300,3 --ipv4-static-route-bfd
10.64.224.16,2.3.6.4,,300,300,3 --max-path 50 --redistribute connected --redistribute
static --graceful-restart-enable --next-hop-recursion-enable true --default-
information-originate-enable true
efa tenant vrf create --tenant "Tnt-5" --name "Tnt-5_VR2" --routing-type "distributed"
--layer3-extension-enable false --local-asn 9001 --rt-type import --rt 103:103 --rt-
type export --rt 103:103 --ipv4-static-route-next-hop 10.64.224.16,9.0.1.0/24,9.1.1.4
--ipv4-static-route-next-hop 10.64.224.16,9.0.2.0/24,9.1.2.4 --max-path 50 --
redistribute connected --redistribute static --graceful-restart-enable --next-hop-
recursion-enable true --default-information-originate-enable true
...
efa tenant create --name "Tnt-6" --type private --vlan-range
411-423,426,1051-1090,1411-1423,1426,1451-1473,1476,1601-1650,2411-2423,2426,2451-2473,
2476,3411-3423,3426,3801-3850 --vrf-count 14 --port
10.64.224.16[0/10:1-4,0/11:1-4,0/12:1,0/19:2-4,0/20:4],10.64.224.17[0/10:1-4,0/11:1-4,0
/12:1,0/19:2-4,0/20:4]
efa tenant po create --tenant "Tnt-6" --name "po613" --description " po613" --min-link-
count 1 --mtu 9000 --negotiation active --number 153 --lacp-timeout short --port
10.64.224.16[0/12:1],10.64.224.17[0/12:1]
...
efa tenant vrf create --tenant "Tnt-6" --name "Tnt-6_VR1" --routing-type "distributed"
--layer3-extension-enable false --rt-type import --rt 102:102 --rt-type export --rt
102:102 --ipv4-static-route-next-hop 10.64.224.17,2.4.1.0/24,2.2.1.3 --ipv4-static-
route-next-hop 10.64.224.17,2.4.1.0/24,2.2.6.4 --ipv6-static-route-bfd
10.64.224.17,2:2:1::3,,300,300,3 --ipv4-static-route-bfd
10.64.224.17,2.2.1.3,,300,300,3 --ipv4-static-route-bfd
10.64.224.16,2.3.6.4,,300,300,3 --ipv4-static-route-bfd
10.64.224.16,2.3.1.3,,300,300,3 --max-path 50 --redistribute connected --redistribute
static --graceful-restart-enable --next-hop-recursion-enable true --default-
information-originate-enable true
...
efa tenant service bgp peer-group create --tenant "Tnt-6" --name "PG_Policy_6_1" --pg-
name 10.64.224.16:pg1:Tnt-6_VR1 --pg-asn 10.64.224.16,pg1,Tnt-6_VR1:1200065535 --pg-
bfd-enable 10.64.224.16,pg1,Tnt-6_VR1:true --pg-bfd
10.64.224.16,pg1,Tnt-6_VR1:450,50,5 --pg-next-hop-self 10.64.224.16,pg1,Tnt-6_VR1:true
--pg-update-source-ip 10.64.224.16,pg1,Tnt-6_VR1:10ab::1 --pg-md5-password
10.64.224.16,pg1,Tnt-6_VR1:<password> --pg-ipv6-uc-nbr-activate
10.64.224.16,pg1,Tnt-6_VR1:true --pg-ipv6-uc-nbr-add-path-capability
10.64.224.16,pg1,Tnt-6_VR1:receive --pg-name 10.64.224.16:pg2:Tnt-6_VR1 --pg-asn
10.64.224.16,pg2,Tnt-6_VR1:1200065535 --pg-bfd-enable 10.64.224.16,pg2,Tnt-6_VR1:true
--pg-bfd 10.64.224.16,pg2,Tnt-6_VR1:450,50,5 --pg-next-hop-self
10.64.224.16,pg2,Tnt-6_VR1:true --pg-update-source-ip
10.64.224.16,pg2,Tnt-6_VR1:10.20.30.1 --pg-md5-password
10.64.224.16,pg2,Tnt-6_VR1:<password> --pg-ipv4-uc-nbr-add-path-capability
10.64.224.16,pg2,Tnt-6_VR1:receive --pg-name 10.64.224.17:pg1:Tnt-6_VR1 --pg-asn
10.64.224.17,pg1,Tnt-6_VR1:1200065535 --pg-bfd-enable 10.64.224.17,pg1,Tnt-6_VR1:true
--pg-bfd 10.64.224.17,pg1,Tnt-6_VR1:450,50,5 --pg-next-hop-self
10.64.224.17,pg1,Tnt-6_VR1:true --pg-update-source-ip
10.64.224.17,pg1,Tnt-6_VR1:10ab::2 --pg-md5-password
10.64.224.17,pg1,Tnt-6_VR1:<password> --pg-ipv6-uc-nbr-activate
10.64.224.17,pg1,Tnt-6_VR1:true --pg-ipv6-uc-nbr-add-path-capability
10.64.224.17,pg1,Tnt-6_VR1:receive --pg-name 10.64.224.17:pg2:Tnt-6_VR1 --pg-asn
10.64.224.17,pg2,Tnt-6_VR1:1200065535 --pg-bfd-enable 10.64.224.17,pg2,Tnt-6_VR1:true
--pg-bfd 10.64.224.17,pg2,Tnt-6_VR1:450,50,5 --pg-next-hop-self
10.64.224.17,pg2,Tnt-6_VR1:true --pg-update-source-ip
10.64.224.17,pg2,Tnt-6_VR1:10.20.30.2 --pg-md5-password
10.64.224.17,pg2,Tnt-6_VR1:<password> --pg-ipv4-uc-nbr-add-path-capability
10.64.224.17,pg2,Tnt-6_VR1:receive
...
efa tenant service bgp peer create --tenant "Tnt-6" --name "customer_Policy_6_1" --
ipv4-uc-nbr 10.64.224.16,Tnt-6_VR1:1.1.1.10,pg2 --ipv4-uc-nbr-bfd
10.64.224.16,Tnt-6_VR1:1.1.1.10,true,50,5001,50 --ipv4-uc-nbr-update-source-ip
10.64.224.16,Tnt-6_VR1:1.1.1.10,10.11.12.13 --ipv4-uc-dyn-nbr
10.64.224.16,Tnt-6_VR1:10.20.50.0/29,pg2,10 --ipv6-uc-nbr
```

```
10.64.224.16,Tnt-6_VR1:110a::10,pg1 --ipv6-uc-nbr-bfd
10.64.224.16,Tnt-6_VR1:110a::10,true,50,5000,50 --ipv6-uc-nbr-update-source-ip
10.64.224.16,Tnt-6_VR1:110a::10,12a::1 --ipv6-uc-dyn-nbr 10.64.224.16,Tnt-6_VR1:13a::/
127,pg1,10 --ipv4-uc-nbr 10.64.224.17,Tnt-6_VR1:1.1.1.11,pg2 --ipv4-uc-nbr-bfd
10.64.224.17,Tnt-6_VR1:1.1.1.11,true,50,5001,50 --ipv4-uc-nbr-update-source-ip
10.64.224.17,Tnt-6_VR1:1.1.1.11,10.11.12.14 --ipv4-uc-dyn-nbr
10.64.224.17,Tnt-6_VR1:10.20.60.0/29,pg2,10 --ipv6-uc-nbr
10.64.224.17,Tnt-6_VR1:110a::11,pg1 --ipv6-uc-nbr-bfd
10.64.224.17,Tnt-6_VR1:110a::11,true,50,5000,50 --ipv6-uc-nbr-update-source-ip
10.64.224.17,Tnt-6_VR1:110a::11,12a::2 --ipv6-uc-dyn-nbr 10.64.224.17,Tnt-6_VR1:23a::/
127,pg1,10
...
efa tenant epg create --tenant "Tnt-6" --name "Tnt-6_Ext_EPG1" --type extension --
switchport-mode trunk --single-homed-bfd-session-type auto --po po613 --vrf Tnt-6_VR1
--ctag-range 1061 --anycast-ip 1061:192.168.189.254/24 --anycast-ipv6
1061:fd01:0:bd:1::fe/64 --local-ip 1061,10.64.224.16:192.168.189.101/24 --local-ipv6
1061,10.64.224.16:fd01:0:bd:1::65/64 --local-ipv6 1061,10.64.224.17:fd01:0:bd:1::66/64
--local-ip 1061,10.64.224.17:192.168.189.102/24 --ctag-description "1061:XCO L3
Extended VLAN" --suppress-arp 1061:false --suppress-nd 1061:false
...
--- Time Elapsed: 7.493054085s ---
(efa:extreme)extreme@node-1:~$
```

5. To run a specified command on one or more devices, run the following device execute CLI command:

```
 efa inventory device execute-cli --ip 10.64.224.17 --command "show version" --config
```

xample:

```
(efa:extreme)extreme@node-1:~$ efa inventory device execute-cli --ip 10.64.224.17 --
command "show version" --config
Execute CLI[success]
+--------------+-----------+---------+--------------+---------+--------
+-------------------------------------------------------------------------------------
----------+
|  IP Address  | Host Name | Fabric  |   Command    | Status  | Reason
|
Output                                          |
+--------------+-----------+---------+--------------+---------+--------
+-------------------------------------------------------------------------------------
----------+
| 10.64.224.17 | Leaf-17   | fabric1 | show version | Success |        |
Leaf-17(config)# OS
ONE                                                                              |
|              |           |         |              |         |        | Copyright (c)
2025 Extreme Networks Inc.                                                         |
|              |           |         |              |         |        | Firmware
Info:                                                                              |
|              |           |         |              |         |        | Base
OS:
   |
|              |           |         |              |         |        | Current
Version:                        10.2.2.0-164
|
|              |           |         |              |         |        | Rollback
Version:                       None                                                |
|              |           |         |              |         |        | Is Current
Version Committed:           True                                                  |
|              |           |         |              |         |        | Switching
Routing (SR):                                                                      |
|              |           |         |              |         |        | Current
Version:                        all-22.2.2.0-164
|
|              |           |         |              |         |        | Rollback
Version:                       None                                                |
```

```
|              |           |         |               |         |        | Is Current
Version Committed:            True                                                  |
|              |           |         |               |         |        | System CPLD
Version:                     Port-CPLD#0:06, Port-CPLD#1:06, Power-CPLD:09         |
|              |           |         |               |         |        | System BIOS
Version:                     0.22.0                                                |
|              |           |         |               |         |        | System BMC
Version:                      1.17                                                 |
|              |           |         |               |         |        | UM-Switch
Primary Version:              1.00.12                                               |
|              |           |         |               |         |        | UM-Switch
Secondary Version:            1.00.9                                                |
|              |           |         |               |         |        | Micro-
controller Version:
1.01.16                                                  |
|              |           |         |               |         |        | Built
Time:
                     |
|              |           |         |               |         |        | Install
Time:
                  |
|              |           |         |               |         |        |
Kernel:                                  6.6.111-yocto-
standard                                |
|              |           |         |               |         |        | System
Uptime:
450h25m58s                                               |
|              |           |         |               |         |        | MicroService
Info:                                                                          |
|              |           |         |               |         |        | BASE
Services:
   |
|              |           |         |               |         |        |
SERVICE                      CURRENT         ROLLBACK      READY       STATE
RESTARTS     |
|              |           |         |               |         |
|                              VERSION
VERSION                                             |
|              |           |         |               |         |        |
-------------------------------------------------------------------------------------
------- |
|              |           |         |               |         |        | api-
gw                       1.1.0           None          true        Running
0            |
|              |           |         |               |         |        | chassis-
mgr                 1.2.0           None          true        Running     0
|
|              |           |         |               |         |        | config-
db                   1.1.0           None          true        Running
0            |
|              |           |         |               |         |        | ingress-
gw                  1.0.0           None          true        Running     0
|
|              |           |         |               |         |        | interface-
mgr               1.2.0           None          true        Running     0            |
|              |           |         |               |         |        | msg-
bus                      1.1.0           None          true        Running
0            |
|              |           |         |               |         |        |
security                     1.1.0           None          true        Running
0            |
|              |           |         |               |         |        | snmp-
agent                  1.1.0           None          true        Running
0            |
```

```
|               |           |         |               |         |        | state-
db                    1.1.0           None          true        Running
0             |
|               |           |         |               |         |        |
ztp                    1.0.0          None          true        Running
0             |
|               |           |         |               |         |        | SR
Services:
     |
|               |           |         |               |         |        |
SERVICE                CURRENT         ROLLBACK      READY       STATE
RESTARTS     |
|               |           |         |               |         |
|                           VERSION
VERSION                                          |
|               |           |         |               |         |        |
-------------------------------------------------------------------------------------
------- |
|               |           |         |               |         |        |
arpnd                  0.1.0          None          true        Running
0             |
|               |           |         |               |         |        |
bfd                    0.1.0          None          true        Running
0             |
|               |           |         |               |         |        |
bgp                    0.1.0          None          true        Running
0             |
|               |           |         |               |         |        |
classifiers            0.1.0          None          true        Running
0             |
|               |           |         |               |         |        | fwd-
hal                  0.1.0          None          true        Running
0             |
|               |           |         |               |         |        |
iah                    0.1.0          None          true        Running
0             |
|               |           |         |               |         |        |
lacp                   1.0.0          None          true        Running
0             |
|               |           |         |               |         |        |
lldp                   1.0.0          None          true        Running
0             |
|               |           |         |               |         |        |
mftm                   0.1.0          None          true        Running
0             |
|               |           |         |               |         |        |
mlag                   0.1.0          None          true        Running
0             |
|               |           |         |               |         |        | sr-api-
gw                 1.0.0          None          true        Running
0             |
|               |           |         |               |         |        |
telegraf               1.0.0          None          true        Running
0             |
|               |           |         |               |         |        |
uftm                   0.1.0          None          true        Running
0             |
|               |           |         |               |         |
|
          |
+--------------+-----------+---------+--------------+---------+--------
+-------------------------------------------------------------------------------------
----------+
```

```
Execute CLI Details
--- Time Elapsed: 7.631168819s ---
```

6. To bulk export the device details in CSV file, run the following REST API:

```
http://127.0.0.1:8082/v1/inventory/switches/inventory-bulk-export
Request Data(Payload) : {"all_device": true}
Response : Inventory_fabric_1732275547.xlsx

(Excel with devices details)
```

7. To reload or reboot a running device, run the following device:

```
efa inventory device reload -ip <deviceIp>
```

Example:

```
efa inventory device reload --ip 10.64.216.16
Warning: This operation will cause the device(s) to reboot and requires all existing
telnet, secure telnets, and SSH sessions to be restarted on the device(s).
----
Warning: EFA access may be lost temporarily if EFA is running on the reloading device.
----
Are you sure you want to reload device(s) 10.64.216.16 (Y/N): y
----
Reload[success]
----
| IP Address   | Status  | Reason |
|--------------|---------|--------|
| 10.64.216.16 | Success |        |
----
Reload Details
Time Elapsed: 49.489569151s
```

8. To view the current discovery interval for a device or fabric, run the following command:

```
efa inventory device discovery-time list
```

Example:

```
efa inventory device discovery-time list --ips 10.64.216.16 ----
| IP Address   | Host Name | Model | Chassis Name | Firmware | Discovery Interval |
Last Discovery Time     | Discovery Reason |
|--------------|-----------|-------|--------------|----------|--------------------|----
--------------------|------------------|
| 10.64.216.16 | ONEOS16     | 3012  | Extreme 8720-32C| 22.1.3.0 | 1 Hour 0 Min
| 2024-11-22 05:23:22 EST | Fabric Configure |
----
Device Discovery Interval Time Details
Time Elapsed: 804.155115ms
```

9. To set the timezone per device or fabric, run the following command:

```
efa inventory device timezone set --<device IP> --timezone <country name>
```

Example:

```
efa inventory device timezone set --ip 10.64.216.44 --timezone "America/Los_Angeles"
+--------------+---------------------+--------+---------+--------+
|  IP Address  |      Timezone       | Fabric | Status  | Reason |
+--------------+---------------------+--------+---------+--------+
| 10.64.216.44 | America/Los_Angeles |        | Success |        |
+--------------+---------------------+--------+---------+--------+
Timezone details
--- Time Elapsed: 14.988099325s ---
```

To set the timezone per device or fabric, run the following command on the SLX device:

```
efa inventory device timezone set --ip 10.64.164.47 --timezone "America/Los_Angeles"
+--------------+---------------------+---------+---------+-------+
| IP Address   | Timezone            | Fabric  | Status  | Reason|
+--------------+---------------------+---------+---------+-------+
| 10.64.164.47 | America/Los_Angeles | fabric2 | Success |       |
```

```
+--------------+--------------------+---------+---------+-------+
Timezone details
--- Time Elapsed: 13.858027763s ---
```

a. To verify the timezone configuration, run the following command:

```
show running-config system clock
system
  clock timezone  America/Los_Angeles
!
```

b. To view the list of timezone per device or fabric, run the following command:

```
efa inventory device timezone list --ip  10.64.216.44
+--------------+---------------------+--------+-------------+
|  IP Address  |      Timezone       | Fabric |  AppState   |
+--------------+---------------------+--------+-------------+
| 10.64.216.44 | America/Los_Angeles |        | cfg-in-sync |
+--------------+---------------------+--------+-------------+
Timezone details
--- Time Elapsed: 30.339321ms ---
```

- To view the list of timezone per device or fabric, run the following command on the SLX device:

```
efa inventory device timezone list --ip 10.64.164.47
+--------------+----------+--------+-------------+
| IP Address   | Timezone | Fabric | AppState    |
+--------------+----------+--------+-------------+
| 10.64.164.47 | Etc/GMT  | fabric2| cfg-in-sync |
+--------------+----------+--------+-------------+
Timezone details
--- Time Elapsed: 66.528502ms ---
```

- Run the following command on the OS ONE devices:

```
SLX#  show running-config clock
clock timezone America/Los_Angeles
```

10. Save the running-config to startup-config on the devices.

- Run the following command on the OS ONEdevices:

```
efa inventory device running-config persist --ip 10.64.164.48
Persist Device(s) Running-Config[success]
+--------------+-----------+--------+---------+--------+
|  IP Address  | Host Name | Fabric | Status  | Reason |
+--------------+-----------+--------+---------+--------+
| 10.64.164.48 | SLX       |        | Success |        |
+--------------+-----------+--------+---------+--------+
Persist Running-Config Details
--- Time Elapsed: 19.216457816s ---
```

- Run the following command on the OS ONE devices:

```
efa inventory device running-config persist --ip 10.64.216.44
Persist Device(s) Running-Config[failed]
+--------------+-----------+--------+--------+----------------------------------+
|  IP Address  | Host Name | Fabric | Status |            Reason                |
+--------------+-----------+--------+--------+----------------------------------+
| 10.64.216.44 | oneos44   |        | Failed | Unsupported on ONEOSSR device for |
|              |           |        |        | ManageSwitchesConfig operation. |
+--------------+-----------+--------+--------+----------------------------------+
Persist Running-Config Details
```

11. To display the list of executions based on the execution ID, limit, or status, run the following command:

```
efa inventory execution show
```

- To display the list of executions based on the limit, run the following command:

```
efa inventory execution show --limit <limit>
```

Example:

```
efa inventory execution show --limit 5
+--------------------------------------
+-----------------------------------------------------------------------------------------
----------+--------+-------------------------+---------------------------
+---------------------------+-----------+
|                  ID
|
URL                                                | METHOD |         STATUS
|        START TIME         |         END TIME          | USER NAME |
+--------------------------------------
+-----------------------------------------------------------------------------------------
----------+--------+-------------------------+---------------------------
+---------------------------+-----------+
| 91f1c1ef-39d5-4efc-9140-b49c09a0fcf0 | /v1/inventory/
switches                                                                              |
GET    | Completed(29.504071ms)  | 2024-12-10T06:26:35-05:00 |
2024-12-10T06:26:35-05:00 |           |
| f27fb42b-5eb4-4171-b38d-6757b23c6d76 | /v1/inventory/switches?
device_ips=10.64.164.47%2C10.64.164.48%2C10.64.184.84%2C10.64.184.85 | DELETE |
Failed(1m0.985896138s)  | 2024-12-10T06:22:08-05:00 | 2024-12-10T06:23:09-05:00
|           |
| 1c1fe29c-934b-432f-a8ea-c23a86af53f8 | /v1/inventory/execution?id=11ba4976-
ba25-416e-b3f7-bcaeb31de33b                              | GET    |
Completed(1.405435234s) | 2024-12-10T06:01:45-05:00 | 2024-12-10T06:01:46-05:00
|           |
| 56b2b7b4-fd6d-4a1f-b6c9-4a5361e24d35 | /v1/inventory/executions?
limit=10&status=failed                                                  | GET    |
Completed(180.236219ms) | 2024-12-10T05:52:33-05:00 | 2024-12-10T05:52:33-05:00
|           |
| 1c2732f2-69b7-4288-aee2-97d47b48ea96 | /v1/inventory/executions?
limit=10&status=all                                                     | GET    |
Completed(27.646312ms)  | 2024-12-10T05:52:28-05:00 | 2024-12-10T05:52:28-05:00
|           |
+--------------------------------------
+-----------------------------------------------------------------------------------------
----------+--------+-------------------------+---------------------------
+---------------------------+-----------+
```

- To display the list of executions based on the status, run the following command:

```
efa inventory execution show --status < failed/succeeded/all>
```

Example:

```
efa inventory execution show --status all
+--------------------------------------
+-----------------------------------------------------------------------------------------
-------+--------+-------------------------+---------------------------
+---------------------------+-----------+
|                  ID                  |
URL                                                | METHOD |         STATUS
|        START TIME         |         END TIME          | USER NAME |
+--------------------------------------
+-----------------------------------------------------------------------------------------
-------+--------+-------------------------+---------------------------
+---------------------------+-----------+
| 49e8eb2b-dbdc-4185-85a1-b929cc93c57c | /v1/inventory/executions?
limit=5&status=all                                                    | GET    |
```

```
Completed(14.200403ms) | 2024-12-10T07:06:50-05:00 | 2024-12-10T07:06:50-05:00
|           |
| 3497ab86-ccc2-43cc-91c8-c865c94430e0 | /v1/inventory/interfaces?
admin_state=all&device_ips=10.64.216.45&oper_state=all&type=all | GET    |
Completed(23.931338ms) | 2024-12-10T07:06:02-05:00 | 2024-12-10T07:06:02-05:00
|           |
| b0ad70f9-aecf-4785-80de-5386c4560a15 | /v1/inventory/portchannels?
device_ips=10.64.216.45                                          | GET    |
Completed(20.241328ms) | 2024-12-10T07:06:02-05:00 | 2024-12-10T07:06:02-05:00
|           |
| 7d9bff1d-49e5-4d9b-bef4-2ed6e19124e1 | /v1/inventory/pim?
device_ip=10.64.216.45                                                | GET    |
Completed(60.901588ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
| 397a110b-9f47-40f5-98c7-a8fc193943b8 | /v1/inventory/mlag/config?
device_ip=10.64.216.44                                            | GET    |
Completed(53.109534ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
| 7082d4e4-98e9-498b-8b18-9aea5c995314 | /v1/inventory/switches?
device_ips=10.64.216.45                                             | GET    |
Completed(29.487787ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
| eb3763db-0c60-4798-bac8-e4f92e350e15 | /v1/inventory/interfaces?
admin_state=all&device_ips=10.64.216.45&oper_state=all&type=All | GET    |
Completed(43.742219ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
| e612e2d6-7806-41b6-9739-39dd35dea5c9 | /v1/inventory/mlag/state?
device_ip=10.64.216.44                                            | GET    |
Completed(80.007139ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
| 10e8c1cb-238f-4c3f-b9af-cf7653123bc4 | /v1/inventory/bgp?
device_ip=10.64.216.45                                                | GET    |
Completed(80.181838ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
| e4989ca0-6d60-4a6d-8ae1-be1b2b164956 | /v1/inventory/mct/management?
device_ip=10.64.216.44                                          | GET    |
Completed(23.455621ms) | 2024-12-10T07:06:01-05:00 | 2024-12-10T07:06:01-05:00
|           |
+--------------------------------------
+-------------------------------------------------------------------------------------
-------+--------+----------------------+-------------------------
+---------------------------+-----------+
--- Time Elapsed: 77.52367ms ---
```

- To display the list of executions based on the execution ID, run the following command:

```
efa inventory execution show --id <ID>
```

Example:

```
efa inventory execution show --id 49e8eb2b-dbdc-4185-85a1-b929cc93c57c
ID        :    49e8eb2b-dbdc-4185-85a1-b929cc93c57c
URL       :    /v1/inventory/executions?limit=5&status=all
Method        :    GET
PARAMETERS    :    {"limit":["5"],"status":["all"]}
STATUS        :    Completed(14.200403ms)
START TIME    :    2024-12-10 07:06:50 -0500 EST
END TIME    :    2024-12-10 07:06:50 -0500 EST
USER        :
LOGS        :
<Logs not available for the Execution ID : 49e8eb2b-dbdc-4185-85a1-b929cc93c57c>
--- Time Elapsed: 244.037614ms ---
```

12. Syslog Server out
13. Certificate

Configure Inventory Device Interface Network Essentials

You can configure device interface network essentials.

About This Task

Follow this procedure to configure device interface network essentials.

Procedure

1. Configure the admin state.
   a. To change the admin state on SLX devices, run the following command:

   For admin state up

```
(efa:user)user@xco-vm-19:~$ efa inventory device interface set-admin-state --ip
10.64.216.36 --if-type eth --if-name 0/1 --state up
+--------------+-----+------+----------------+--------------+---------+--------+
|   DeviceIP   | ID  | Name | Interface Type | Admin Status | Result  | Reason |
+--------------+-----+------+----------------+--------------+---------+--------+
| 10.64.216.36 | 193 | 0/1  | ethernet       | up           | Success |        |
+--------------+-----+------+----------------+--------------+---------+--------+
Interface Details
--- Time Elapsed: 6.587336949s ---
(efa:user)user@xco-vm-19:~
```

   For admin state down

```
(efa:user)user@xco-vm-19:~$ efa inventory device interface set-admin-state --ip
10.64.216.36 --if-type eth --if-name 0/1 --state down
+--------------+-----+------+----------------+--------------+---------+--------+
|   DeviceIP   | ID  | Name | Interface Type | Admin Status | Result  | Reason |
+--------------+-----+------+----------------+--------------+---------+--------+
| 10.64.216.36 | 193 | 0/1  | ethernet       | down         | Success |        |
+--------------+-----+------+----------------+--------------+---------+--------+
Interface Details
--- Time Elapsed: 6.587336949s ---
(efa:user)user@xco-vm-19:~
```

   b. To change the admin state on OS ONE devices, run the following command:

   For admin state up

```
ONEOSL6# show running-config interface ethernet 0/1
interface ethernet 0/1
no shutdown
!
```

   For admin state down

```
ONEOSL6# show running-config interface ethernet 0/1
interface ethernet 0/1
shutdown
!
```

2. Configure MTU.
   a. To configure MTU on SLX device, run the following command:

   To set MTU

```
slx160.49(conf-if-eth-0/54)# do show running-config
interface ethernet 0/54
interface Ethernet 0/54
mtu 4000
no shutdown
!
```

   To unset MTU

```
slx160.49(conf-if-eth-0/54)# do show running-config interface ethernet 0/54
interface Ethernet 0/54
```

```
ip mtu 4000
no shutdown
!
```

b. To configure MTU on OS ONE devices, run the following command:

To set MTU

```
ONEOSL6(config-if-eth-0/1)# show running-config interface ethernet 0/1
interface ethernet 0/1
mtu 4000
no shutdown
!
```

To unset MTU

```
ONEOSL6(config-if-eth-0/1)# show running-config interface ethernet 0/1
interface ethernet 0/1
no shutdown
!
```

3. Configure IP MTU.

a. To configure IP MTU on SLX devices, run the following command:

To set IP MTU

```
slx160.49(conf-if-eth-0/54)# do show running-config interface ethernet 0/54
interface Ethernet 0/54
ip mtu 4000
no shutdown
!
```

To unset IP MTU

```
slx160.49(conf-if-eth-0/54)# do show running-config interface ethernet 0/54
interface Ethernet 0/54
no shutdown
!
```

b. To configure IP MTU on OS ONE devices, run the following command:

To set IP MTU

```
ONEOSL6(config-if-eth-0/1)# show running-config interface ethernet 0/1
interface ethernet 0/1
ip mtu 4000
no shutdown
!
```

To unset IP MTU

```
ONEOSL6(config-if-eth-0/1)# show running-config interface ethernet 0/1
interface ethernet 0/1
no shutdown
!
```

4. Configure breakout ports.

You can break a port into multiple interfaces, such as breaking one 40G port into four 10G ports.

a. To configure breakout ports on SLX devices, run the following command:

Set breakout port

```
slx160.49(config-connector-0/50)# interface ethernet 0/50
slx160.49(conf-if-eth-0/50)# shutdown
slx160.49(conf-if-eth-0/50)# hardware
slx160.49(config-hardware)# connector 0/50
slx160.49(config-connector-0/50)# breakout mode 4x10g
slx160.49(config-connector-0/50)# do show running-config interface ethernet 0/50:1-4
interface Ethernet 0/50:1
shutdown
!
```

```
interface Ethernet 0/50:2
shutdown
!
interface Ethernet 0/50:3
shutdown
!
interface Ethernet 0/50:4
shutdown
!
```

Unset breakout port

```
slx160.49(config-connector-0/50) # no breakout
slx160.49(config-connector-0/50)# do show running-config interface ethernet 0/50:1-4
% No entries found.
slx160.49(config-connector-0/50)# do show running-config interface ethernet 0/50
interface Ethernet 0/50
shutdown
!
```

b. To configure breakout ports on OS ONE devices, run the following command:

Set breakout port

```
32c(config-connector-0/1)# breakout 4x10g
32c(config-connector-0/1)# do sh running-config interface ethernet 0/1
32c(config-connector-0/1)#
32c(config-connector-0/1)#
32c(config-connector-0/1)# do sh running-config interface ethernet 0/1:1-4
interface ethernet 0/1:1
shutdown
!
interface ethernet 0/1:2
shutdown
!
interface ethernet 0/1:3
shutdown
!
interface ethernet 0/1:4
shutdown
!
```

Unset breakout port

```
32c(config-connector-0/1)# no breakout
32c(config-connector-0/1)#
32c(config-connector-0/1)#
32c(config-connector-0/1)# do sh running-config interface ethernet 0/1:1-4
Error: Invalid interface ethernet 0/1:1
Error: Invalid interface ethernet 0/1:2
Error: Invalid interface ethernet 0/1:3
Error: Invalid interface ethernet 0/1:4
32c(config-connector-0/1)# do sh running-config interface ethernet 0/1
interface ethernet 0/1
shutdown
!
```

5. Set the port speed on a device.

a. To set the port speed on SLX device, run the following command

```
1000 1Gbps
10000 10Gbps
100000 100Gbps
25000 25Gbps
40000 40Gbps
auto Auto Detection (default)
```

b. To set the port speed on OS ONE device, run the following command

```
ONEOSL6(config-if-eth-0/2)# show running-config interface ethernet 0/2
interface ethernet 0/2
speed 1000
no shutdown
!
```

6. Configure the forwarding error correction (FEC).
   a. To set the FEC on the SLX device, run the following command

```
auto-negotiation
disabled
fc-fec
rs-fec
```

   b. To set the FEC on the OS ONE device, run the following command

```
ONEOSL6(config-if-eth-0/1)# show running-config interface ethernet 0/1
interface ethernet 0/1
fec rs-fec
no shutdown
!
```

Configure Device Level Network Essentials

You can configure device-level network essentials.

About This Task
Follow this procedure to configure device-level network essentials.

Procedure

1. Show device configuration.

```
sh running-config system  ntp
system
  ntp
    server 10.64.220.47
    enable
  !
!
```

2. Configure an NTP server.
   a. Run the **`efa inventory device ntp server create`** command.

```
 efa inventory device ntp server create [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
      --ntp-ip string           Server ip
      --auth-key int            authentication key id. Valid values are 1-65535.
Not supported on ExtremeOneOS devices
      --auth-key-name string    authentication key name. Not supported on
ExtremeOneOS devices
      --encryption-type string  Encryption type. Valid values are md5, sha1. Not
supported on ExtremeOneOS devices
```

```
      --trusted-key              Trusted key (optional)
      --fabric string            Fabric name
```

Note
The parameters `auth-key`, `auth-key-name`, and `encryption-type` is not supported on Extreme OS ONE devices.

```
efa inventory device  ntp server create --ip 10.64.216.36 --ntp-ip 10.64.220.47
+--------------+--------------+----------+---------------+-----------------
+-------------+--------+---------+--------+
|  IP Address  |  Server IP   | Auth Key | Auth
Key Name | Encryption Type | Trusted Key | Fabric | Status  | Reason |
+--------------+--------------+----------+---------------+-----------------
+-------------+--------+---------+--------+
| 10.64.216.36 | 10.64.220.47 | 0        |
|                 | false       |         | Success |        |
+--------------+--------------+----------+---------------+-----------------
+-------------+--------+---------+--------+
Ntp server details
--- Time Elapsed: 6.599601397s ---
(efa:user)user@xco-vm-19:~$
```

b. Enable or disable an NTP server.

To turn off an NTP server on a device:

```
efa inventory device ntp disable-server --enable yes --ip 10.20.246.10
```

To turn on an NTP server on a device:

```
efa inventory device ntp disable-server --enable no --ip 10.20.246.10
```

To turn off an NTP server at the fabric level:

```
efa inventory device ntp disable-server --enable yes --fabric clos_fabric
```

To turn on an NTP server at the fabric level:

```
efa inventory device ntp disable-server --enable no --fabric clos_fabric
```

c. Delete an NTP server.

```
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
```

d. List the NTP servers configured using XCO.

```
efa inventory device ntp server list --ip 10.20.246.10
efa inventory device ntp server list --fabric clos_fabric
```

3. Configure an SNMP server.

a. Create an SNMP community and SNMP group.

```
efa inventory device snmp community create --ip 10.139.44.153 --name community1
--group group1 --enable-read-access --view view1
```

b. Delete an SNMP community.

```
efa inventory device snmp community delete --name tempv2community
--ip 10.20.246.10
```

c. Lists SNMP communities

```
efa inventory device snmp community list --ip 10.20.246.1-2
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| IP Address  | Community Name              | Group   | Read view
| Write view | Notify view | AppState        |
+-------------+-----------------------------+---------+-----------+------------
```

```
+-------------+-----------------+
| 10.20.246.2 | $9$jcpLTpu9FlqFFcWN7uhHig== |         |           |
|             | cfg-not-managed |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| 10.20.246.2 | $9$smklvisSghOZEQvXJKBDeA== | cgrp1   | efav3View
| efav3View  |             | cfg-in-sync     |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| 10.20.246.1 | $9$Y4Cy6IcflH9B0KM932JUNw== | v2group | efav3View
| efav3View  |             | cfg-not-managed |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| 10.20.246.1 | $9$yS/jSs6dBJyG5o0yy+coKw== |         |           |
|             | cfg-not-managed |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| 10.20.246.1 | $9$kR+plantb3TRIgkiBkV83Q== | v2group | efav3View
| efav3View  |             | cfg-not-managed |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| 10.20.246.1 | $9$mdDWwNPg2OmO4Fqnodl+Bw== | gtest1  | efav3View
| efav3View  |             | cfg-not-managed |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
| 10.20.246.1 | $9$smklvisSghOZEQvXJKBDeA== | cgrp1   | efav3View
| efav3View  |             | cfg-in-sync     |
+-------------+-----------------------------+---------+-----------+------------
+-------------+-----------------+
Snmp community details
```

d. Creates an SNMP v2c or v3 host

e. Delete an SNMP v2c or v3 host.

```
efa inventory device snmp host delete
--host-ip 20.20.20.2 --community tempv2community --ip 10.20.246.10

efa inventory device snmp host delete
--host-ip 30.30.30.2 --user v3user --ip 10.20.246.10
```

f. Create an SNMP user and SNMP group

```
efa inventory device snmp user create --ip 10.139.44.153 --name user1
--group group1 --view view1
```

g. Delete an SNMP user.

```
efa inventory device snmp user delete --name tempv3user --ip 10.20.246.10
```

h. List SNMP users.

```
efa inventory device snmp user list --ip 10.139.44.153
```

i. Create an SNMP View.

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 detail
```

For example, see *ExtremeCloud Orchestrator Command Reference Guide*.

Configure Device Level Settings

You can configure device-level settings.

About This Task

Follow this procedure to configure device-level settings.

Note

OS ONE supports `maint-mode-enable` and `maint-mode-enable-on-reboot`. Note that `maint-mode-enable` is now privileged and doesn't persist after reboot.

Procedure

1. Configure an MCT Delay Timer.

   The following example specifies a delay of 250 seconds before an MCT cluster bring up:

```
efa inventory device setting update  --mct-bring-up-delay 250 --ip 10.64.216.36-37
+--------------+--------------------+---------+-------+-------+
|  IP ADDRESS  |        NAME        | STATUS  | VALUE | ERROR |
+--------------+--------------------+---------+-------+-------+
| 10.64.216.36 | MCT Bring-up Delay | Success | 250   |       |
+--------------+--------------------+---------+-------+-------+
| 10.64.216.37 | MCT Bring-up Delay | Success | 250   |       |
+--------------+--------------------+---------+-------+-------+
--- Time Elapsed: 113.358263ms ---
```

2. Verify the MCT Delay Timer configuration on SLX and OS ONE devices.

SLX running-config	OS ONE running-config

`cluster fs-cluster-1`
`  bringup-delay 250`
	`mlag`
	`  bringup-delay 250`
	`!`

3. Perform a device health check and backup configuration.

   By default, health checks are disabled on SLX and OS ONE devices when they are registered. To enable them, use the **`health-check-enable`**, **`health-check-interval`**, and **`health-check-heartbeat-miss-threshold`** commands. Additionally, you can configure XCO to regularly back up device configurations at a set interval (default is every 6 minutes).

```
efa inventory device setting update --ip "<device_ip> --health-check-enable
<Yes>                --health-check-interval <6m> --health-check-heartbeat-miss-
threshold <2>
```

   - Enable device health check.

```
efa inventory device setting update --ip 10.x.x.x --health-check-enable yes
```

   - Configure health check interval.

```
efa inventory device setting update --ip 10.x.x.x --health-check-interval 30mins
```

   - Configure health check threshold.

```
efa inventory device setting update --ip 10.x.x.x --health-check-heartbeat-
missthreshold 2
```

   - View device health status.

```
./efa inventory device health status --ip 10.x.x.x
+--------------------------+----------+
| NAME                     | VALUE    |
+--------------------------+----------+
| health check status      | enabled  |
```

```
+--------------------------+----------+
| heartbeat miss counter   | 0        |
+--------------------------+----------+
| heartbeat miss threshold | 2        |
+--------------------------+----------+
| Health Check Interval    | 6m       |
+--------------------------+----------+
| system maintenance mode  | disabled |
+--------------------------+----------+
| drift and reconcile UUID |          |
+--------------------------+----------+
```

- (Optional) Disable device health check.

  ```
  efa inventory device setting update --ip 10.x.x.x --health-check-enable no
  ```

4. Configure RMA on devices.

   Ensure that you have completed the device level settings first before configuring RMA and DRC.

   To enable RMA on supported devices, ensure the following prerequisites are met:

   - **Periodic configuration backup:** Enable regular configuration backups on all devices that may require RMA to ensure the latest configuration file is available for recovery.
   - **Maintenance mode on reboot:** Enable maintenance mode on all devices upon reboot.

5. View the drift reconcile history.

   ```
   efa inventory drift-reconcile history --device-ip <ip-addr>
   ```

   Example:

   ```
   ./efa inventory drift-reconcile history

   +--------------------------------------+--------------+---------+--------+---------------------
   +------------------------------+
   | UUID                                 | Device IP    | Status  | Reason | Type
   | Start Time                     |
   +--------------------------------------+--------------+---------+---------+---------------------
   +------------------------------+
   | c1b9cf42-8c51-409b-8e97-ad0add08ea76 | 10.64.216.16 | success | manual | drift-and-
   reconcile | 2024-12-02 06:40:22 -0500 EST |
   +--------------------------------------+--------------+---------+--------+---------------------
   +------------------------------+
   Drift Reconcile
   ```

   - Reconcile: False

     ```
     efa inventory drift-reconcile execute --ip <device_ip>
     ```

     Example:

     ```
     ./efa inventory drift-reconcile execute --ip 10.64.216.16
     +--------------+-----------+----------------------------------------+---------
     +--------+
     | IP ADDRESS   | RECONCILE | UUID
     | STATUS  | REASON |
     +--------------+-----------+----------------------------------------+---------
     +--------+
     | 10.64.216.16 | No        | ce3196a6-752a-4084-8f01-53bea7dd296f
     | Success |        |
     +--------------+-----------+----------------------------------------+---------
     +--------+
     Drift and Reconcile Execute
     Execute the CLI to get details : efa inventory drift-
     reconcile detail --uuid ce3196a6-752a-4084-8f01-53bea7dd296f

     efa inventory drift-reconcile detail --uuid <UUId>
     ```

Example:

```
./efa inventory drift-reconcile detail --uuid ce3196a6-752a-4084-8f01-53bea7dd296f
+----------------------------------+----------------------------------------+
| NAME                             | VALUE                                  |
+----------------------------------+----------------------------------------+
| UUID                             | ce3196a6-752a-4084-8f01-53bea7dd296f |
+----------------------------------+----------------------------------------+
| Device IP                        | 10.64.216.16                           |
+----------------------------------+----------------------------------------+
| Status                           | success                                |
+----------------------------------+----------------------------------------+
| Execution Reason                 | manual                                 |
+----------------------------------+----------------------------------------+
| operation                        | drift-only                             |
+----------------------------------+----------------------------------------+
| Inventory Status                 | inventory-dr-success                   |
+----------------------------------+----------------------------------------+
| Is Inventory config Refreshed    | true                                   |
+----------------------------------+----------------------------------------+
| Inventory Duration               | 39.222935ms                            |
+----------------------------------+----------------------------------------+
| Fabric Status                    | fabric-dr-success                      |
+----------------------------------+----------------------------------------+
| Is Fabric config Refreshed       | false                                  |
+----------------------------------+----------------------------------------+
| Fabric Duration                  | 70.28409ms                             |
+----------------------------------+----------------------------------------+
| Policy Status                    | policy-dr-success                      |
+----------------------------------+----------------------------------------+
| Is Policy config Refreshed       | false                                  |
+----------------------------------+----------------------------------------+
| Policy Duration                  | 185.059726ms                           |
+----------------------------------+----------------------------------------+
| Tenant Status                    | tenant-dr-success                      |
+----------------------------------+----------------------------------------+
| Is Tenant config Refreshed       | false                                  |
+----------------------------------+----------------------------------------+
| Tenant Duration                  | 30.729824ms                            |
+----------------------------------+----------------------------------------+
| Device Update Count              | 2                                      |
+----------------------------------+----------------------------------------+
| Device Update Total Duration     | 1m40.031198775s                        |
+----------------------------------+----------------------------------------+
| Maintenance Mode Disable         |                                        |
| Duration                         |                                        |
+----------------------------------+----------------------------------------+
| Start Time                       | 2024-12-02 07:45:27 -0500 EST          |
+----------------------------------+----------------------------------------+
| Last Modified                    | 2024-12-02 07:47:33 -0500 EST          |
+----------------------------------+----------------------------------------+
| Duration                         | 2m6.008938557s                         |
+----------------------------------+----------------------------------------+

Inventory Service Response:
Config Drift: Interface Config
+------+---------------+------------------------------+
| NAME | APP STATE     | CHILD CONFIG                 |
+------+---------------+------------------------------+
| 0/3 | cfg-refreshed  | Interface Config Admin State |
+------+---------------+------------------------------+
Fabric Service Response:
Policy Service Response:
Tenant service Response:
```

- Reconcile: True

```
efa inventory drift-reconcile execute --ip <device_ip> --reconcile
```

Example:

```
./efa inventory drift-reconcile execute --ip 10.64.216.16 --reconcile
+--------------+-----------+--------------------------------------+---------
+--------+
| IP ADDRESS   | RECONCILE | UUID
| STATUS  | REASON |
+--------------+-----------+--------------------------------------+---------
+--------+
| 10.64.216.16 | Yes       | 9821ccc3-e599-4660-b454-8cac440fde57
| Success |        |
+--------------+-----------+--------------------------------------+---------
+--------+
Drift and Reconcile Execute
Execute the CLI to get details : efa inventory drift-
reconcile detail --uuid 9821ccc3-e599-4660-b454-8cac440fde57
```

```
efa inventory drift-reconcile detail --uuid
```

Example:

```
./efa inventory drift-reconcile detail --uuid 9821ccc3-e599-4660-b454-8cac440fde57
+--------------------------------+--------------------------------------+
| NAME                           | VALUE                                |
+--------------------------------+--------------------------------------+
| UUID                           | 9821ccc3-e599-4660-b454-8cac440fde57 |
+--------------------------------+--------------------------------------+
| Device IP                      | 10.64.216.16                         |
+--------------------------------+--------------------------------------+
| Status                         | success                              |
+--------------------------------+--------------------------------------+
| Execution Reason               | manual                               |
+--------------------------------+--------------------------------------+
| operation                      | drift-and-reconcile                  |
+--------------------------------+--------------------------------------+
| Inventory Status               | inventory-dr-success                 |
+--------------------------------+--------------------------------------+
| Is Inventory config Refreshed | true                                 |
+--------------------------------+--------------------------------------+
| Inventory Duration             | 13.461259467s                        |
+--------------------------------+--------------------------------------+
| Fabric Status                  | fabric-dr-success                    |
+--------------------------------+--------------------------------------+
| Is Fabric config Refreshed     | false                                |
+--------------------------------+--------------------------------------+
| Fabric Duration                | 105.883705ms                         |
+--------------------------------+--------------------------------------+
| Policy Status                  | policy-dr-success                    |
+--------------------------------+--------------------------------------+
| Is Policy config Refreshed     | false                                |
+--------------------------------+--------------------------------------+
| Policy Duration                | 13.161342866s                        |
+--------------------------------+--------------------------------------+
| Tenant Status                  | tenant-dr-success                    |
+--------------------------------+--------------------------------------+
| Is Tenant config Refreshed     | false                                |
+--------------------------------+--------------------------------------+
| Tenant Duration                | 47.973437ms                          |
+--------------------------------+--------------------------------------+
| Device Update Count            | 2                                    |
+--------------------------------+--------------------------------------+
| Device Update Total Duration   | 1m42.363570046s                      |
+--------------------------------+--------------------------------------+
| Maintenance Mode Disable       |                                      |
```

```
| Duration                      |                                       |
+-------------------------------+---------------------------------------+
| Start Time                    | 2024-12-02 07:53:08 -0500 EST         |
+-------------------------------+---------------------------------------+
| Last Modified                 | 2024-12-02 07:55:56 -0500 EST         |
+-------------------------------+---------------------------------------+
| Duration                      | 2m48.30651742s                        |
+-------------------------------+---------------------------------------+

Inventory Service Response:
Config Drift: Interface Config
+------+---------------+------------------------------+
| NAME | APP STATE     | CHILD CONFIG                 |
+------+---------------+------------------------------+
| 0/3  | cfg-refreshed | Interface Config Admin State |
+------+---------------+------------------------------+

Reconcile Status:
+-------------------+---------------+---------------+
| CONFIG-TYPE       | APP STATE     | ERROR-MESSAGE |
+-------------------+---------------+---------------+
| SnmpUser          | Not-Attempted |               |
| SnmpHost          | Not-Attempted |               |
| SnmpUseVrf        | Not-Attempted |               |
| DeviceTimezone    | Not-Attempted |               |
| DeviceSetting     | Not-Attempted |               |
| ThresholdMonitor  | Not-Attempted |               |
| BreakoutInterface | Not-Attempted |               |
| InterfaceConfig   | Success       |               |
| SnmpGroup         | Not-Attempted |               |
| SnmpCommunity     | Not-Attempted |               |
| MMOnReboot        | Not-Attempted |               |
| NtpDisable        | Not-Attempted |               |
| NtpServer         | Not-Attempted |               |
| SecureSetting     | Not-Attempted |               |
| NtpAuthKey        | Not-Attempted |               |
| SnmpView          | Not-Attempted |               |
+-------------------+---------------+---------------+
Fabric Service Response:
Policy Service Response:
Tenant service Response:
--- Time Elapsed: 34.425612ms ---
```

6. Bring the device admin up or down.
   a. To bring the device admin up, run the following command:

```
.efa inventory admin-state up --ip 10.64.216.36
AdminStateUp [success]
```

   Run the following command on OS ONE:

```
Show running-config

interface ethernet 0/2
no shutdown
!
```

   b. To bring the device admin down, run the following command:

```
efa inventory admin-state down --ip 10.64.216.36
AdminStateUp [success]
```

   Run the following command on OS ONE devices:

```
Show running-config

interface ethernet 0/2
shutdown
!
```

7. Run the config tracking command on SLX and OS ONE devices:

   The config-tracking is always enabled on OS ONE devices.

   - On SLX devices

```
SLX#  show config-drift-track values
Config Drift Timestamp: 1733221373
Config Drift Counter: 3057

SLX# show config-drift-track status
Config Drift Tracking is OFF

SLX# show config-drift-track status
Config Drift Tracking is ON
```

   - On OS ONE devices

```
device# sh system internal cdb timestamp last-updated
Version    Updated time                          User              Address
========   ===================================   ===============   =============
3168       2026-02-04 16:31:03.06484427 +0000 UTC   admin          10.64.193.251

device#
```

8. Enable prefix-independent-convergence on the SLX device.

```
efa inventory device setting update --ip 10.20.24.18 --prefix-independent-
convergencestatic yes
+-------------+--------------------------+---------+---------+-------+
| IP ADDRESS  | NAME                     | STATUS  | VALUE   | ERROR |
+-------------+--------------------------+---------+---------+-------+
| 10.20.24.18 | Static Prefix Independent | Success | Yes     |       |
|             | Convergence Enabled       |         |         |       |
+-------------+--------------------------+---------+---------+-------+
Warning: The best practice is to clear all routes after configuring Static PIC. You can
also reload the device.
Execute the CLI to clear : efa inventory device clear route-all --ip 10.20.24.18
Execute the CLI to reload : efa inventory device reload --ip 10.20.24.18
```

9. Run the following command on SLX and OS ONE devices to enable BGP PIC:

   - On SLX devices

```
efa inventory device setting update --ip 10.20.24.10 --prefix-independent-
convergence
Yes
Warning: The best practice is to clear all routes after configuring PIC. You can
also
reload the device.
Execute the CLI to clear : efa inventory device clear route-all --ip 10.20.24.1
```

   - On OS ONE devices

```
router bgp
    local-as 4200000000
    router-id 172.31.254.137
   as-notation as-plain
    address-family ipv4 unicast
      graceful-restart
      prefix-independent-convergence
      use-multiple-paths ibgp maximum-paths 8
      use-multiple-paths ebgp maximum-paths 8
      network 172.31.254.226/32
      network 172.31.254.137/32
      activate
    !
    address-family l2vpn evpn
      graceful-restart
      activate
```

```
    !
    address-family ipv6 unicast
      prefix-independent-convergence
      activate
    !
```

10. Configure certificate expiry.

    Configure certificate expiry on the SLX devices.

```
efa inventory device  setting update --ip 10.64.216.36 --crypto-cert-expiry-info
80 --crypto-cert-expiry-minor 70 --crypto-cert-expiry-major 50 --crypto-cert-expiry-
critical 40
+--------------+--------------------------------+---------+-------+-------+
|  IP ADDRESS  |              NAME              | STATUS  | VALUE | ERROR |
+--------------+--------------------------------+---------+-------+-------+
| 10.64.216.36 | Crypto Certificate Expiry for  | Success | 40    |       |
|              | Critical Level                 |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Crypto Certificate Expiry For  | Success | 80    |       |
|              | Info Level                     |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Crypto Certificate Expiry for  | Success | 50    |       |
|              | Major Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Crypto Certificate Expiry for  | Success | 70    |       |
|              | Minor Level                    |         |       |       |
+--------------+--------------------------------+---------+-------+-------+
--- Time Elapsed: 127.681505ms ---
```

    Configure certificate expiry on the OS ONE devices.

```
device# sh running-config system
system
  hostname node1
  grub
    username root password-hashed
grub.pbkdf2.sha512.10000.EB7341398786C524C8E378ADB079E479F767644461ABDBCE63C7FB4E95CE89
FA382FD4E9EC46DC800646A31E7C81CC8492224F159E91B5EA
C524628D45C4CF6C.ACB55BA6F710C5910DFCDA035827291EF51D4A49DDF686803948D09C80F697C6164543
E9EBD94E3AA8021DCB6CDE4B9D1BF7AF59AD2F2469E965D513E1A58993
  !
  aaa
    authentication
      admin-user admin password-hashed $6$EU8FrfKa2uv2UxOK$Vj/GIs/Z7YEt0ecjhctGLXWh6v/
sKrnKwnWlwezMvqhrTlojDB6CSd/MVEfYbGIst93h91dgP9PTuwViczpXS0
      password-attributes
        max-password-age 90
        expiry-alert
          info 45
          minor 20
          major 10
          critical 3
        !
      !
    !
    token-validator xco_jwt_token_validator
      ssl-profile-id xco_jwt_ssl_profile_id
    !
  !
  global
    mtu 9216
    ip mtu 9100
    ipv4 anycast-gateway-mac 02:01:01:01:01:01
    ipv6 anycast-gateway-mac 02:01:01:01:01:02
    ecmp-max-path 32
```

```
    mac-address-table aging-time 1800
    l2 mac-move-detection enable
    l2 mac-move-detection threshold 20
  !
  grpc-server xco_grpc_server_443
    certificate-id xco_grpc_ssl_profile_id
    port 443
    enable
  !
  certificate-manager
    expiry-alert
      critical 5
      major 15
      minor 30
      info 60
    !
  !
  ssh-server DEFAULT
    mac  hmac-sha2-256 hmac-sha2-512
    key-exchange  curve25519-sha256 diffie-hellman-group-exchange-sha256 diffie-
hellman-group16-sha512 diffie-hellman-group18-sha512 diffie-hellman-group14-sha256
    key rsa 4096
    vrf mgmt-vrf
    enable
  !
  snmp-server DEFAULT
    user efav3User auth sha auth-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b priv aes priv-password-hashed
0x5c5eaf97a540da6fb23a01835fc721cc87a5de2b
    vrf mgmt-vrf
      host 10.64.208.102 efav3User
        version 3
        udp-port 162
      !
      enable
    !
  !
  logging
    remote-server 10.64.208.102
      secure-forwarding tls
      mode-transport tcp
      vrf mgmt-vrf
      tls-profile-id xco_ca_ssl_profile_id
    !
  !
  maintenance-mode
    enable-on-reload
  !
!
```

Update Extreme OS ONE device settings in xco inventory.

```
efa inventory device setting update --fabric fab1 --maint-mode-enable Yes --maint-mode-
enable-on-reboot Yes --mct-bring-up-delay 100 --health-check-enable Yes --health-check-
interval 15m --health-check-heartbeat-miss-threshold 3 --config-backup-periodic-enable
Yes --config-backup-interval 120m --number-of-config-backups 5 --prefix-independent-
convergence Yes --maximum-load-sharing-paths 32 --crypto-cert-expiry-info 60 --crypto-
cert-expiry-minor 30 --crypto-cert-expiry-major 15 --crypto-cert-expiry-critical 5
--password-expiry-info 45 --password-expiry-minor 20 --password-expiry-major 10 --
password-expiry-critical 3
+--------------+--------------------------------+---------+-------+-------+
|  IP ADDRESS  |              NAME              | STATUS  | VALUE | ERROR |
+--------------+--------------------------------+---------+-------+-------+
| 10.64.208.25 | Maintenance Mode Enable On     | Success | Yes   |       |
```

```
|               | Reboot                          |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Password Expiry for Major       | Success | 10    |       |
|               | Level                           |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Route Load-Sharing Maximum      | Success | 32    |       |
|               | Paths                           |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Password Expiry for Critical    | Success | 3     |       |
|               | Level                           |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Password Expiry For Info Level  | Success | 45    |       |
+               +---------------------------------+---------+-------+-------+
|               | Config Backup Interval          | Success | 120m  |       |
+               +---------------------------------+---------+-------+-------+
|               | Health Check Interval           | Success | 15m   |       |
+               +---------------------------------+---------+-------+-------+
|               | Password Expiry for Minor       | Success | 20    |       |
|               | Level                           |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Config Backup Max Count         | Success | 5     |       |
+               +---------------------------------+---------+-------+-------+
|               | Periodic Config Backup Enabled  | Success | Yes   |       |
+               +---------------------------------+---------+-------+-------+
|               | Crypto Certificate Expiry for   | Success | 5     |       |
|               | Critical Level                  |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Crypto Certificate Expiry For   | Success | 60    |       |
|               | Info Level                      |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Health Check Enabled            | Success | Yes   |       |
+               +---------------------------------+---------+-------+-------+
|               | MCT Bring-up Delay              | Success | 100   |       |
+               +---------------------------------+---------+-------+-------+
|               | Prefix Independent Convergence  | Success | Yes   |       |
|               | Enabled                         |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Crypto Certificate Expiry for   | Success | 15    |       |
|               | Major Level                     |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Crypto Certificate Expiry for   | Success | 30    |       |
|               | Minor Level                     |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Health Check Heartbeat Miss     | Success | 3     |       |
|               | Threshold                       |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Maintenance Mode Enable         | Success | Yes   |       |
+---------------+---------------------------------+---------+-------+-------+
| 10.64.208.24  | Crypto Certificate Expiry for   | Success | 5     |       |
|               | Critical Level                  |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Crypto Certificate Expiry For   | Success | 60    |       |
|               | Info Level                      |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Health Check Heartbeat Miss     | Success | 3     |       |
|               | Threshold                       |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Password Expiry for Critical    | Success | 3     |       |
|               | Level                           |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Password Expiry for Major       | Success | 10    |       |
|               | Level                           |         |       |       |
+               +---------------------------------+---------+-------+-------+
|               | Config Backup Interval          | Success | 120m  |       |
```

```
+              +--------------------------------+---------+-------+-------+
|              | Maintenance Mode Enable        | Success | Yes   |       |
+              +--------------------------------+---------+-------+-------+
|              | MCT Bring-up Delay             | Success | 100   |       |
+              +--------------------------------+---------+-------+-------+
|              | Maintenance Mode Enable On     | Success | Yes   |       |
|              | Reboot                         |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Crypto Certificate Expiry for  | Success | 15    |       |
|              | Major Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Health Check Enabled           | Success | Yes   |       |
+              +--------------------------------+---------+-------+-------+
|              | Health Check Interval          | Success | 15m   |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry For Info Level | Success | 45    |       |
+              +--------------------------------+---------+-------+-------+
|              | Password Expiry for Minor      | Success | 20    |       |
|              | Level                          |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Crypto Certificate Expiry for  | Success | 30    |       |
|              | Minor Level                    |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Route Load-Sharing Maximum     | Success | 32    |       |
|              | Paths                          |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Prefix Independent Convergence | Success | Yes   |       |
|              | Enabled                        |         |       |       |
+              +--------------------------------+---------+-------+-------+
|              | Config Backup Max Count        | Success | 5     |       |
+              +--------------------------------+---------+-------+-------+
|              | Periodic Config Backup Enabled | Success | Yes   |       |
+--------------+--------------------------------+---------+-------+-------+
Warning: Route load-sharing maximum paths configuration will not take effect until
reloaded.
Execute the CLI to reload : efa inventory device reload --fabric fab1
Warning: It is recommended to clear all routes after configuring PIC, reloading device
may also be done instead.
Execute the CLI to clear : efa inventory device clear route-all --fabric fab1
Execute the CLI to reload : efa inventory device reload --fabric fab1
--- Time Elapsed: 32.425693899s ---
```

11. Verify the telnet server configuration on SLX and OS ONE devices.

    On SLX devices

    ```
    telnet server use-vrf mgmt-vrf shutdown
    ```

    On OS ONE devices

    ```
    system
      ssh-server DEFAULT
        vrf mgmt-vrf
        enable
      !
    ```

12. Verify the cipher configuration on SLX and OS ONE devices.

    On SLX devices

    ```
    ssh server cipher aes128-cbc
    ```

    On OS ONE devices

    ```
    system
      ssh-server DEFAULT
        cipher  aes128-cbc
      !
    !
    ```

The following table describes the possible values by platform type:

SLX	OS ONE	gNMI
aes128-ctr	aes128-ctr	AES128_CTR
aes192-ctr	aes192-ctr	AES192_CTR
aes256-ctr	aes256-ctr	AES256_CTR
aes128-cbc	aes128-cbc	AES128_CBC
aes192-cbc	aes192-cbc	AES192_CBC
aes256-cbc	aes256-cbc	AES256_CBC
aes128-gcm@openssh.com	aes128-gcm-openssh	AES128_GCM
aes256-gcm@openssh.com	aes256-gcm-openssh	AES256_GCM
chacha20-poly1305@openssh.com	chacha20-poly1305-openssh	CHACHA20_POLY1305
blowfish-cbc		
cast128-cbc		
arcfour		
arcfour128		
arcfour256		
rijndael-cbc@lysator.liu.se		
3des-cbc		
non-cbc		

13. Verify the MAC algorithm configuration on SLX and OS ONE devices.

On SLX devices

```
ssh server mac hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,umac-128-
etm@openssh.com,hmac-sha2-256,hmac-sha2-512,hmac-sha1
```

On OS ONE devices

```
system
  ssh-server DEFAULT
    mac  hmac-sha1 hmac-sha1-96 hmac-sha2-256 hmac-sha2-512-etm-openssh
  !
!
```

The following table describes the possible values by platform type:

SLX	OS ONE	OS ONE gNMI
hmac-sha2-256	hmac-sha2-256	HMAC_SHA2_256
hmac-sha2-512	hmac-sha2-512	HMAC_SHA2_512
hmac-sha1	hmac-sha1	HMAC_SHA1
hmac-sha1-96	hmac-sha1-96	HMAC_SHA1_96
hmac-sha2-256-etm@openssh.com	hmac-sha2-256-etm-openssh	HMAC_SHA2_256_ETM
hmac-sha2-512-etm@openssh.com	hmac-sha2-512-etm-openssh	HMAC_SHA2_512_ETM

SLX	OS ONE	OS ONE gNMI
hmac-sha1-etm@openssh.com	hmac-sha1-etm-openssh	HMAC_SHA1_ETM
hmac-md5		
hmac-md5-96		
hmac-ripemd160		
hmac-ripemd160@openssh.com		
umac-64@openssh.com		
umac-128@openssh.com		
hmac-sha1-96-etm@openssh.com		
hmac-md5-etm@openssh.com		
hmac-ripemd160-etm@openssh.com		
umac-64-etm@openssh.com		
umac-128-etm@openssh.com		
hmac-ripemd160-etm@openssh.com		

14. Verify the key exchange algorithm configuration on SLX and OS ONE devices.

    On SLX devices

    ```
    ssh server key-exchange curve25519-sha256,curve25519-sha256@libssh.org,diffie-hellman-group-exchange-sha256,diffie-hellman-group16-sha512,diffie-hellman-group18-sha512
    ```

    On OS ONE devices

    ```
    system
      ssh-server DEFAULT
        key-exchange  diffie-hellman-group-exchange-sha1 diffie-hellman-group16-sha512
    diffie-hellman-group14-sha256 diffie-hellman-group-exchange-sha256
      !
    !
    ```

    The following table describes the possible values by platform type:

SLX	OS ONE	gNMI
curve25519-sha256	curve25519-sha256	CURVE25519_SHA256
ecdh-sha2-nistp256	ecdh-sha2-nistp256	ECHD_SHA2_NISTP256
ecdh-sha2-nistp384	ecdh-sha2-nistp384	ECHD_SHA2_NISTP384
ecdh-sha2-nistp521	ecdh-sha2-nistp521	ECHD_SHA2_NISTP521
diffie-hellman-group-exchange-sha256	diffie-hellman-group-exchange-sha256	DIFFIE_HELLMAN_GROUP_KEYEXCHANGE_SHA256
diffie-hellman-group16-sha512	diffie-hellman-group16-sha512	DIFFIE_HELLMAN_GROUP16_SHA512

SLX	OS ONE	gNMI
diffie-hellman-group18-sha512	diffie-hellman-group18-sha512	DIFFIE_HELLMAN_GROUP18_SHA512
diffie-hellman-group14-sha256	diffie-hellman-group14-sha256	DIFFIE_HELLMAN_GROUP14_SHA256
diffie-hellman-group14-sha1	diffie-hellman-group14-sha1	DIFFIE_HELLMAN_GROUP14_SHA1
diffie-hellman-group-exchange-sha1	diffie-hellman-group-exchange-sha1	DIFFIE_HELLMAN_GROUP_KEYEXCHANGE_SHA1
curve25519-sha256@libssh.org	curve25519-sha256-libssh	CURVE25519_SHA256_LIB_SSH
diffie-hellman-group1-sha1		

15. Verify the force default password change configuration on SLX and OS ONE devices.

    On SLX devices

    ```
    password-attributes force-default-password-change
    ```

    On OS ONE devices

    Note
    For Extreme OS ONE, the **`--force-default-password-change`** attribute is always enabled. Even if you set it to disable, it still shows as Enable in **`efa show-run-config`** command output.

    ```
    efa inventory device  secure settings update --force-default-password-change disable
    secure settings update[Success]
    — Time Elapsed: 201.536554ms —
    (efa:extreme)extreme@tpvm44:~$ efa inventory device secure settings apply --ip
    10.64.216.44
    ----------------------------

     IP Address   Status   Warning
    ----------------------------

    10.64.216.44   Success
    ----------------------------
    Secure settings apply
    — Time Elapsed: 11.091146011s —
    (efa:extreme)extreme@tpvm44:~$ efa show-running-config | grep "secure se"
    efa inventory device secure settings enable
    efa inventory device secure settings update --min-tls-version 1.2 --mac-
    algorithm hmac-sha1,hmac-sha1-96,hmac-sha2-256 --cipher aes128-ctr,aes128-cbc,aes128-
    gcm-openssh,aes256-ctr --key-exchange-algorithm curve25519-sha256,curve25519-sha256-
    libssh --telnet Enable --max-password-age 50 --force-default-password-change Enable
    efa inventory device secure settings apply --ip "10.64.216.44"
    efa inventory device secure settings enable
    efa inventory device secure settings update --min-tls-version 1.2 --mac-
    algorithm hmac-sha1,hmac-sha1-96,hmac-sha2-256 --cipher aes128-ctr,aes128-cbc,aes128-
    gcm-openssh,aes256-ctr --key-exchange-algorithm curve25519-sha256,curve25519-sha256-
    libssh --telnet Enable --max-password-age 50 --force-default-password-change Enable
    efa inventory device secure settings apply --ip "10.64.216.45"
    timed out waiting for input: auto-logout
    ```

16. Verify the maximum password age configuration on SLX and OS ONE devices.

    On SLX devices

    ```
    password-attributes max-password-age 300
    ```

On OS ONE devices

```
system
  aaa
    authentication
      password-attributes
        max-password-age 300
       !
     !
  !
!
```

Note
On OS ONE:

- The **force-default-password-change** command is enabled by default - It cannot be disabled, even if you try to set it to **disable** in global secure settings.
- XCO shows it as disabled in global settings, but it remains enabled at the device level.

Configure IP Fabric BGP-EVPN Single Rack Non-Clos LCM

You can configure a BGP-EVPN single rack Non-Clos LCM in an IP fabric.

About This Task
Follow this procedure to configure a BGP-EVPN single rack Non-Clos LCM in an IP fabric.

Procedure

1. Create a fabric.

```
efa fabric create --name <fabric-name> --type <clos | non-clos> --description
<description>
```

2. Display the fabric settings.

```
efa fabric setting show --name fs --advanced
+--------------------------------+-------------------------+
|              NAME              |          VALUE          |
+--------------------------------+-------------------------+
| Fabric Name                    | fs                      |
+--------------------------------+-------------------------+
| Link IP Range                  | 10.12.5.0/20            |
+--------------------------------+-------------------------+
| Loopback Scheme                | granular                |
+--------------------------------+-------------------------+
| Loopback IP Range              | 172.31.254.0/24         |
+--------------------------------+-------------------------+
| RouterID Loopback IP Range     | 172.32.128.20/24        |
+--------------------------------+-------------------------+
| VTEP Loopback IP Range         | 172.31.66.20/24         |
+--------------------------------+-------------------------+
| Loopback Port Number           | 1                       |
+--------------------------------+-------------------------+
| VTEP Loopback Port Number      | 2                       |
+--------------------------------+-------------------------+
| Rack ASN Block                 | 4200000100-4200055534   |
+--------------------------------+-------------------------+
| Rack Border Leaf ASN Block     | 4200065580-4200065620   |
+--------------------------------+-------------------------+
```

```
| Single Rack Deployment          | No                     |
+---------------------------------+------------------------+
| P2P IP Type                     | numbered               |
+---------------------------------+------------------------+
| Any cast MAC                    | aabb.aabb.aabb         |
+---------------------------------+------------------------+
| IPV6 Any cast MAC               | aa00.aa00.aa00         |
+---------------------------------+------------------------+
| MAC Aging Timeout               | 60                     |
+---------------------------------+------------------------+
| MAC Aging Conversational        | 80                     |
| Timeout                         |                        |
+---------------------------------+------------------------+
| MAC Move Limit                  | 5                      |
+---------------------------------+------------------------+
| Duplicate MAC Timer             | 10                     |
+---------------------------------+------------------------+
| Dampening Delay Timer           | 6                      |
+---------------------------------+------------------------+
| Duplicate MAC Timer MAX Count   | 8                      |
+---------------------------------+------------------------+
| BFD Enable                      | Yes                    |
+---------------------------------+------------------------+
| BFD Tx                          | 400                    |
+---------------------------------+------------------------+
| BFD Rx                          | 450                    |
+---------------------------------+------------------------+
| BFD Multiplier                  | 4                      |
+---------------------------------+------------------------+
| MaxPaths                        | 6                      |
+---------------------------------+------------------------+
| AllowAsIn                       | 0                      |
+---------------------------------+------------------------+
| MTU                             | 1500                   |
+---------------------------------+------------------------+
| IPMTU                           | 1300                   |
+---------------------------------+------------------------+
| MCT Link IP Range               | 10.15.30.10/20         |
+---------------------------------+------------------------+
| MCT PortChannel                 | 62                     |
+---------------------------------+------------------------+
| LACP Timeout                    | short                  |
+---------------------------------+------------------------+
| Control Vlan                    | 4090                   |
+---------------------------------+------------------------+
| Control VE                      | 4090                   |
+---------------------------------+------------------------+
| Rack Underlay EBGP Peer Group   | underlay               |
+---------------------------------+------------------------+
| Rack Overlay EBGP Peer Group    | overlay                |
+---------------------------------+------------------------+
| BGP MultiHop                    | 6                      |
+---------------------------------+------------------------+
| Rack MCT Scheme                 | dynamic                |
+---------------------------------+------------------------+
| Rack MCT Ports                  |                        |
+---------------------------------+------------------------+
| Rack Low Density MCT Ports      |                        |
+---------------------------------+------------------------+
| Configure Overlay Gateway       | Yes                    |
+---------------------------------+------------------------+
| Overlay Gateway Broadcast       | No                     |
| Local Bias Enable               |                        |
+---------------------------------+------------------------+
```

```
| VNI Auto Map                    | No                      |
+---------------------------------+-------------------------+
| Backup Routing Enable           | Yes                     |
+---------------------------------+-------------------------+
| Backup Routing IPv4 Range       | 10.40.50.0/24           |
+---------------------------------+-------------------------+
| Backup Routing IPv6 Range       | fd40:4040:4040:2::/120  |
+---------------------------------+-------------------------+
| MD5 Password Enable             | Yes                     |
+---------------------------------+-------------------------+
| MD5 Password                    | *****                   |
+---------------------------------+-------------------------+
| BGP Dynamic Peer Listen Limit   | 100                     |
+---------------------------------+-------------------------+
| VNI Domain Auto Offset          | 0                       |
+---------------------------------+-------------------------+
--- Time Elapsed: 80.056364ms ---
```

To update the fabric settings, run the **`efa fabric setting update`** command.

```
efa fabric setting  update  ? | grep EXTREME-OS-ONE
      --mac-aging-timeout string                          MAC Aging Timeout. SLX device
<NUMBER: 0|60-86400>, EXTREME-OS-ONE device <NUMBER 0|60-38400>
      --mac-aging-conversation-timeout string             MAC Conversational Aging time in
seconds<NUMBER: 0|60-100000>,  Not Supported for EXTREME-OS-ONE devices
      --mac-move-limit string                             MAC move detect limit, SLX device
<NUMBER: 5-500>, EXTREME-OS-ONE device <NUMBER: 3-500>
      --overlay-gateway-broadcast-local-bias-enable string   Enable or Disable Overlay Gateway
Broadcast Local Bias on Multi-Homed Leaf and Border Leaf devices <STRING Yes/No default No, Not
```

Supported for Extreme OS ONE devices

```
--bfd-tx string                                            BFD desired min transmit
interval in milliseconds SLX device <NUMBER: 50-30000>, EXTREME-OS-ONE device <NUMBER:
10-30000>
      --bfd-rx string                                      BFD desired min receive
interval in milliseconds SLX device <NUMBER: 50-30000>, EXTREME-OS-ONE device <NUMBER:
10-30000>
      --optimized-replication-enable string                Enable Optimized
replication <STRING Yes/No>, Not Supported for EXTREME-OS-ONE devices
      --mdtgroup-range string                              IPv4 multicast address
range in IP prefix format. Example: 239.0.0.0/8, Not Supported for EXTREME-OS-ONE
devices
      --default-mdtgroup string                            Default IPv4 multicast
address. Address must be from the MDT group range, Not Supported for EXTREME-OS-ONE
devices
      --bgp-dynamic-peer-listen-limit string               Denotes the maximum
number of dynamic bgp peers that can be operational across the VRFs in the SLX
<NUMBER: 1-2400, default 100>, Not Supported for EXTREME-OS-ONE devices
(efa:user)user@xco-vm-19:~$
```

To add one device to the fabric, run the **`efa fabric device add`** command.

To add multiple devices to the fabric, run the **`efa fabric device add-bulk`** command.

```
efa fabric device add --name fb2 --ip 10.64.216.36 --username <username> --password <password>
Inventory Device(s) Registration[success]
+----+--------------+-----------+-------+--------------+----------------------
+-------------------------------------------------------------------+
| ID |  IP Address  | Host Name | Model | Chassis
Name |       Firmware       |                                                                   |
+----+--------------+-----------+-------+--------------+----------------------
+-------------------------------------------------------------------+
```

```
| 7  | 10.64.216.36 | TB1-TOS36 | 3012  | Extreme
8720-32C| OSOneSR-22.2.2.0-070 | Device with IP: 10.64.216.36 already registered in the          |
|    |              |           |       |
|                      | application                                                          |
+----+--------------+-----------+-------+--------------+----------------------
+-------------------------------------------------------------+
Device Details
Add Device(s) [Success]

    Addition of Leaf device with ip-address = 10.64.216.36 [Succeeded]
Validate Fabric [Failed]
    Missing Links
    Rack r1, should have 2 number of devices
Error :     fabric validation failed

Add device to fabric

Usage:
  efa fabric device add [flags]

Flags:
      --name string            Name of the fabric
      --ip stringArray         Pair of Device IPs for rack, single IP for non-rack
      --role string            Device Role (leaf|spine|super-spine|border-leaf)
      --leaf-type string       Leaf Type (single-homed | multi-homed)
      --hostname string        Host Name
      --asn string             ASN
      --vtep-loopback string   VTEP Loopback
      --loopback string        Loopback Port Number
      --pod string             Name of the pod
      --username string        Username for the device
      --password string        Password for the device (default "<prompt>")
      --rack stringArray       Rack name
  -h, --help                   help for add
--- Time Elapsed: 13.378139094s ---

Note: the error in above output is expected as only one device for the fabric is added.
```

3. Configure the fabric.

```
efa fabric show --name fb2

Fabric Name: fb2, Fabric Description: , Fabric Type: non-clos, Fabric Status:
configure-success, Fabric Health: Green
+--------------+------+-----------+------------+------+--------------+-------------
+--------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK | HOST NAME |    ASN     | ROLE | DEVICE
STATE |  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+------+-----------+------------+------+--------------+-------------
+--------------------+-----------------+---------+-------+
| 10.64.216.37 | r1   | TB1-TOS37 | 4200000000 | leaf | provisioned
| cfg in-sync | NA               | NA              | 2       | 1     |
| 10.64.216.36 | r1   | TB1-TOS36 | 4200000000 | 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, BFDU/BFDD
- BFD Update/Delete, VRFU - VRF Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn
Update, SYS - System Properties Update
```

```
MD5 - BGP MD5 Password, BGPU/BGPD - Router BGP Update/Delete, BGPLL
- BGP Listen Limit, POU - Port Channel Update, BGPAFD/BGPAFU/BGPAFC - BGP Address
family Delete/Update/Create, L2VPNC/L2VPNU/L2VPND - BGP L2VPN Address family Create/
Update/Delete , BGPPGU/BGPPGD - BGP Peer Group Update/Delete, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties,
INTIP - Interface IP, BGP - Router BGP, VRFM - VRF Member, BGPPG - BGP Peer Group
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: 35.961813ms ---
```

a. Complete the global configuration.

```
protocol lacp
  lacp system-priority 32768
!mac-address-table aging-time 1801
mtu 9215
ip mtu 9101
ip anycast-gateway-mac 02:01:01: 01:01:01
ipv6 anycast-gateway-mac 02:01:01: 01:01:02

bfd
  profile <fabric_name>
    interval 300 min-rx 350 multiplier 3
  !
  profile mlag
    interval 300 min-rx 350 multiplier 3
  !
  profile default
     interval 300 min-rx 300 multiplier 3
   !
```

b. Display the physical topology configurations.

```
efa fabric topology show physical --name clos_fabric
+--------------+-----------+---------------------------+------------------
| SOURCE       | SOURCE    | DESTINATION | DESTINATION | SOURCE
|  NODE IP     | NODE ROLE | NODE IP     | NODE ROLE   | NODE INTERFACE
+--------------+-----------+-------------+-------------+----------------
| 10.24.80.135 | LEAF      |10.25.225.163 | LEAF       | 0/39
| 10.24.80.135 | LEAF      |10.25.225.163 | LEAF       | 0/40
| 10.24.80.135 | LEAF      |10.24.80.134  | LEAF       | 0/46
| 10.24.80.134 | LEAF      |10.24.80.135  | LEAF       | 0/46


+----------------+----------------+-------
| DESTINATION    | SOURCE DEVICE  | DESTINATION DEVICE
| NODE INTERFACE | MULTI HOMED    | MULTI HOMED
+----------------+----------------+-------
| 0/39           | TRUE           | FALSE
| 0/40           | TRUE           | FALSE
| 0/46           | TRUE           | TRUE
| 0/46           | TRUE           | TRUE
```

Configure Ethernet interface.

```
interface ethernet 0/27
  description clusterPeerIntfMember
  channel-group 63 mode active
  lacp rate normal
  no shutdown
!
interface ethernet 0/28
```

```
  description clusterPeerIntfMember
  channel-group 63 mode active
  lacp rate normal
  no shutdown
!
interface ethernet 0/31
  description clusterPeerIntfMember
  channel-group 63 mode active
  lacp rate normal
  no shutdown
!
interface ethernet 0/32
  description clusterPeerIntfMember
  channel-group 63 mode active
  lacp rate normal
  no shutdown
```

Configure MCT port-channel.

```
interface port-channel 63
  description MCTPeerInterface
  ip address 10.20.20.2/31
  no shutdown
!
```

.

4. Configure overlay topology.
   a. Verify the EVPN configuration.

```
evpn-instance efa_fabric
    nve efa_fabric
    ignore-as
    duplicate-host-detection
      mac-detection
        move-count 3
        detection-time 5
        dampening-time 5
      !
    !
  !
```

   b. Verify the following configurations on SLX and OS ONE devices.

SLX configuration	OS ONE configuration	

evpn efa_fabric
 route-target both
auto ignore-as
 rd auto
 duplicate-mac-timer 5
max-count 3
!
	evpn-instance
	efa_fabric
	  nve efa_fabric
	  ignore-as
	  duplicate-host-
	detection
	  mac-detection
	  move-count 3
	  detection-time 5
	  dampening-time 5
	  !
	  !
	!
		For OS ONE, the evpn-instance is configured in the VRF.

c. Complete the following overlay configurations.

Note
XCO 4.0.2 does not support Auto VNI Offset configuration for Brownfield Fabric Service overlay networks. Configuring an offset under overlay networks causes VNI allocation mismatches between XCO and OS ONE devices.

```
overlay-networks
  default-nve base
  default-vni-offset 0
  nvo base
    member nve base
    member vni-domain efa_fabric
    split-horizon base
  !
  nvo efa_fabric
    member nve efa_fabric
    member vni-domain efa_fabric
    split-horizon efa_fabric
  !
  nve base
  nve efa_fabric
    source-interface loopback 2
  !
  vni-domain base
  vni-domain efa_fabric
    auto-offset 10
  !
!
```

Verify the following configuration on SLX and OS ONE devices:

SLX configurations	OS ONE configurations

`overlay-gateway efa_fabric`
`  ip interface Loopback 2`
`  map vni auto`
`  activate`
	`overlay networks`
	`  nvo efa_fabric`
	`  member nve efa_fabric`
	`  member vni-domain efa_fabric`
	`  split-horizon efa_fabric`
	`  nve efa_fabric`
	`source-interface loopback 2`
	`  vni-domain efa_fabric`

5. Remove the devices from the fabric.

```
efa fabric device remove --ip 10.64.184.87 --name efa_fabric
Remove Device(s) [Success]
Removal of  device with ip-address = 10.64.184.87 [Succeeded]

--- Time Elapsed: 25.0356082s ---
```

When a device is removed from a fabric, the underlay and overlay configurations are removed from the device and the device is no longer a member of the fabric.

If the `no-device-cleanup` option is used, the configuration pushed by the automation engine is not cleaned up from the fabric devices. Removal of a device from the fabric does not delete the device from inventory. You must explicitly delete the device from inventory.

The following example removes five devices and does not clean up the configuration.

```
efa fabric device remove --ip 10.24.48.131,10.24.51.135,10.25.225.58,
10.24.51.131,10.24.80.139 --name BLR_FABRIC --no-device-cleanup
```

6. Delete the fabric.

   The following example deletes the `efa_fabric` fabric:

```
efa fabric delete --name efa_fabric  --force
Delete Fabric [Success]
```

   Deletion of a fabric is not allowed if the fabric has one or more devices. You must delete all the devices from the fabric prior to deleting the fabric. Forced deletion of a fabric removes the devices from the fabric but not from the inventory.

7. Debug the fabric.

   a. To identify drift in device configuration, run the following command:

```
efa fabric debug device drift --name fs --device-ip 10.64.188.58 --reconcile
Fabric Service Response:
Config Drift: Global Config
+---------------------+-------------+----------------+
|       CONFIG        |  APP STATE  | EXPECTED VALUE |
+---------------------+-------------+----------------+
| Mtu                 | cfg-in-sync | 1500           |
| IPMtu               | cfg-in-sync | 1300           |
| AnycastMac          | cfg-in-sync | aabb.aabb.aabb |
| IPV6AnycastMac      | cfg-in-sync | aa00.aa00.aa00 |
| ProtocolLacpEnabled | cfg-in-sync | true           |
| LacpSystemPriority  | cfg-in-sync | 32768          |
+---------------------+-------------+----------------+
Config Drift: EVPN
+------+---------------+---------------------------+
| NAME |   APP STATE   |       CHILD CONFIG        |
+------+---------------+---------------------------+
| fs   | cfg-refreshed | SwEvpnName                |
| fs   | cfg-refreshed | RouteTarget               |
| fs   | cfg-refreshed | DuplicateMacTimerMaxCount |
| fs   | cfg-refreshed | DuplicateMacTimer         |
| fs   | cfg-refreshed | DampeningTime             |
| fs   | cfg-refreshed | OverlayNve                |
+------+---------------+---------------------------+
Config Drift: Overlay Gateway
+------+-------------+------------------------+
| NAME |  APP STATE  |      CHILD CONFIG      |
+------+-------------+------------------------+
| fs   | cfg-in-sync | VtepLoopbackPortNumber |
| fs   | cfg-in-sync | NvoName                |
| fs   | cfg-in-sync | NvoMemberNve           |
| fs   | cfg-in-sync | NvoMemberVniDomain     |
| fs   | cfg-in-sync | NvoMemberSplitHorizon  |
| fs   | cfg-in-sync | NveName                |
| fs   | cfg-in-sync | VniDomainName          |
| fs   | cfg-in-sync | VniDomainAutoOffset    |
+------+-------------+------------------------+
Config Drift: Cluster
+---------+---------------+----------------------------------+
|  NAME   |   APP STATE   |           CHILD CONFIG           |
+---------+---------------+----------------------------------+
| Cluster | cfg-refreshed | ClusterName                      |
| Cluster | cfg-refreshed | MCTPeerName::0:Port-channel:62   |
| Cluster | cfg-refreshed | MCTPeerKeepalivePrimaryName::0   |
| Cluster | cfg-refreshed | MCTPeerKeepaliveSecondaryName::0 |
| Cluster | cfg-refreshed | MCTPeerKeepalivePrimaryType::0   |
| Cluster | cfg-refreshed | MCTPeerKeepaliveSecondaryType::0 |
```

```
+---------+---------------+---------------------------------+
Config Drift: Interface
+------+-------------+----------+---------------------------------+
| NAME |  APP STATE  | INT TYPE |          CHILD CONFIG           |
+------+-------------+----------+---------------------------------+
| 1    | cfg-in-sync | loopback | IP:1:loopback:172.32.128.165/32 |
| 2    | cfg-in-sync | loopback | IP:2:loopback:172.31.66.35/32   |
+------+-------------+----------+---------------------------------+
Config Drift: Router BGP
+--------+---------------+---------------------------------+
|  TYPE  |   APP STATE   |          CHILD CONFIG           |
+--------+---------------+---------------------------------+
| Global | cfg-refreshed | PeerGroupInfo                   |
| Global | cfg-in-sync   | BgpNeighbor                     |
| Global | cfg-refreshed | BgpMCTNeighbor                  |
| Global | cfg-refreshed | RouterID                        |
| Global | cfg-refreshed | LocalAsn                        |
| Global | cfg-in-sync   | BfdMultiplier                   |
| Global | cfg-in-sync   | BfdTx                           |
| Global | cfg-in-sync   | BfdRx                           |
| Global | cfg-refreshed | BgpIPV4Network                  |
| Global | cfg-refreshed | BgpIPV4NetworkGracefulRestart   |
| Global | cfg-refreshed | BgpL2EVPNNetworkGracefulRestart |
| Global | cfg-in-sync   | BgpL2EVPNRetainRtAll            |
| Global | cfg-refreshed | BgpIPV4NetworkMaxPath           |
+--------+---------------+---------------------------------+
Config Drift: Vrf
+-------------+-------------+---------------+
|    NAME     |  APP STATE  | CHILD CONFIG  |
+-------------+-------------+---------------+
| default-vrf | cfg-in-sync | VrfInterfaces |
+-------------+-------------+---------------+
+-------------+---------+-------+
| CONFIG TYPE | STATUS  | ERROR |
+-------------+---------+-------+
| MCT         | Success |       |
| routerbgp   | Success |       |
| evpn        | Success |       |
+-------------+---------+-------+
--- Time Elapsed: 40.821901898s ---
```

The key difference between the **`efa fabric debug device drift`** and **`efa inventory drift-reconcile`** commands is as follows:

- **`efa fabric debug device drift`**: Displays or reconciles configuration drift between a device and the intended fabric configuration.
- **`efa inventory drift-reconcile`**: Detects configuration drift on a device and performs the reconciliation process.

b. To check the configuration generation reason for a particular fabric device in a fabric, run the following command:

```
efa fabric debug config-gen-reason --name fs --device 10.64.216.16
+-------------+-------------+----------------+----------------+------------
+-------------------+--------------------+-----------------------
+-----------------------+
|  IP ADDRESS  | DEVICE ROLE |
INTERFACE TYPE | INTERFACE NAME | LINK STATE | REMOTE
IP ADDRESS | REMOTE DEVICE ROLE | REMOTE INTERFACE TYPE | REMOTE INTERFACE NAME |
+-------------+-------------+----------------+----------------+------------
+-------------------+--------------------+-----------------------
+-----------------------+
| 10.64.216.16 | Leaf        | ethernet
```

```
| 0/15           | deleted    | 10.64.216.14
| Spine              | ethernet              | 0/4                   |
+--------------+-------------+----------------+----------------+------------
+-------------------+--------------------+-----------------------
+-----------------------+
config generate reason [Success]
--- Time Elapsed: 61.410027ms ---
(efa:user)user@user:~$ efa fabric debug config-
gen-reason --name fs --device 10.64.216.14
+--------------+-------------+----------------+----------------+------------
+-------------------+--------------------+-----------------------
+-----------------------+
|  IP ADDRESS  | DEVICE ROLE
| INTERFACE TYPE | INTERFACE NAME | LINK STATE | REMOTE
IP ADDRESS | REMOTE DEVICE ROLE | REMOTE INTERFACE TYPE | REMOTE INTERFACE NAME |
+--------------+-------------+----------------+----------------+------------
+-------------------+--------------------+-----------------------
+-----------------------+
| 10.64.216.14 | Spine       | ethernet
| 0/4            | deleted    | 10.64.216.16
| Leaf               | ethernet              | 0/15                  |
+--------------+-------------+----------------+----------------+------------
+-------------------+--------------------+-----------------------
+-----------------------+
config generate reason [Success]
```

c. To clear the underlay or overlay configuration from the device and recovers the device from erroneous conditions, run the following command:

```
efa fabric debug clear-config --device 10.64.184.87
Clear Config [Success]
```

d. To list all the lock statuses for fabric services, run the following command:

```
efa fabric debug service lock
+---------------------+--------+--------+
|      Lock type      | Locked | Reason |
+---------------------+--------+--------+
| REST Lock           | false  |        |
+---------------------+--------+--------+
| Service Lock        | false  |        |
+---------------------+--------+--------+
| DB Lock             | false  |        |
+---------------------+--------+--------+
| Device 10.64.188.58 | false  |        |
+---------------------+--------+--------+
| Device 10.64.188.61 | false  |        |
+---------------------+--------+--------+
| Device 10.64.188.60 | false  |        |
+---------------------+--------+--------+
| Device 10.64.188.59 | false  |        |
+---------------------+--------+--------+
Lock Status
--- Time Elapsed: 22.217329ms ---
```

e. To enables debug of smartdrc for fabric services, run the following command:

```
efa fabric debug service smartdrc
Enable or Disable SmartDrc Config Succeeded, ConfiugredVal : false

--- Time Elapsed: 26.737011ms ---

efa fabric debug service smartdrc --enable
Enable or Disable SmartDrc Config Succeeded, ConfiugredVal : true
```

The Smart Drift and Reconciliation automatically detect, analyse, and fix discrepancies between the target network configuration and the actual running configuration on

network devices (for example, switches and routers). It ensures network stability, security, reduced outage, automated recovery, and compliance by maintaining a single source of truth and removing the need for manual, error-prone configuration fixes.

f. To compute the fabric health, run the following command:

```
efa fabric debug health compute --name fs
Health Computation for fabric fs is success
--- Time Elapsed: 246.193363ms ---
```

8. Display the health of the fabric.

```
efa fabric health --name fs
=============================================================================
Fabric Name                          :  fs
Fabric Type                          :  non-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 |
+--------------+------+---------------------+-------------------+---------------+
| 10.64.184.86 | Leaf | Green               | Green             | Green         |
| 10.64.184.87 | Leaf | Green               | Green             | Green         |
+--------------+------+---------------------+-------------------+---------------

(efa:user)user@user:~/omkar/efa$ efa fabric health show --name fs
=============================================================================================
Fabric Name                            :  fs
Fabric Type                            :  non-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 |
+--------------+------+---------------------+-------------------+---------------+
| 10.64.188.58 | Leaf | Green               | Red               | Red           |
| 10.64.188.59 | Leaf | Green               | Red               | Red           |
+--------------+------+---------------------+-------------------+---------------+

=============================================================================
--- Time Elapsed: 47.594465ms ---
(efa:user)user@user:~/omkar/efa$ efa fabric health show --name fs --detail
=============================================================================
Fabric Name                            :  fs
Fabric Type                            :  non-clos
Fabric Health                          :  Red
 Fabric Status                         :  configure-success
 Fabric Level Physical Topology Health   :  Green
-----------------------------------------------------------------------------
Fabric Device Health

 Device IP [Role]                                 :  10.64.188.58 [Leaf]
  Device Health                                   :  Red
  Configuration State Health                      :  Green
    Dev State                                     :  provisioned
    App State                                     :  cfg in-sync
  Operational State Health                        :  Red
    Cluster Health                                :  Red
      Operational State                           :  true
      Peer Operational State                      :  true
      Peer Keepalive Primary  Operational State   :  true
      Peer Keepalive Secondary Operational State  :  false
    Physical Topology Device Health               :  Green
```

```
    Underlay Topology Device Health             :  Green
 ------------------------------------------------------
 Device IP [Role]                               :  10.64.188.59 [Leaf]
  Device Health                                 :  Red
  Configuration State Health                    :  Green
    Dev State                                   :  provisioned
    App State                                   :  cfg in-sync
  Operational State Health                      :  Red
    Cluster Health                              :  Red
      Operational State                         :  true
      Peer Operational State                    :  true
      Peer Keepalive Primary  Operational State  :  true
      Peer Keepalive Secondary Operational State :  false
    Physical Topology Device Health             :  Green
    Underlay Topology Device Health             :  Green
-------------------------------------------------------------------------
```

9. Display the devices in the IP fabric which is in error state.

```
$ efa fabric error show
+-------------+--------------+-------+------------+--------------------------------+
| FABRIC NAME |  IP ADDRESS  | ROLE  | ERROR TYPE |          ERROR REASON          |
+-------------+--------------+-------+------------+--------------------------------+
| fab-non-clos| NA           | NA    | clos error | No Devices                     |
| fab-clos    | 10.20.246.12 | Leaf  | clos error | Leaf Device 10.20.246.12 not   |
|             |              |       |            | connected to Spine Device      |
|             |              |       |            | 10.20.246.19                   |
| fab-clos    | 10.20.246.12 | Leaf  | clos error | Leaf Device 10.20.246.12 not   |
|             |              |       |            | connected to Spine Device      |
|             |              |       |            | 10.20.246.20                   |
| fab-clos    | 10.20.246.19 | Spine | clos error | Spine Device 10.20.246.19      |
|             |              |       |            | not connected to Leaf Device   |
|             |              |       |            | 10.20.246.12                   |
| fab-clos    | 10.20.246.20 | Spine | clos error | Spine Device 10.20.246.20      |
|             |              |       |            | not connected to Leaf Device   |
|             |              |       |            | 10.20.246.12                   |
| fab-clos    | 10.20.246.19 | Spine | clos error | Spine Device 10.20.246.19      |
|             |              |       |            | connected to Spine Device      |
|             |              |       |            | 10.20.246.20                   |
| fab-clos    | 10.20.246.20 | Spine | clos error | Spine Device 10.20.246.20      |
|             |              |       |            | connected to Spine Device      |
|             |              |       |            | 10.20.246.19                   |
+-------------+--------------+-------+------------+--------------------------------+
```

10. To display the list of executions, run the fabric execution command.

```
efa fabric execution show
+--------------------------------------+--------------------------------
+------------------------+----------------------+----------------------+-----------+
|                  ID                  |            COMMAND
|         STATUS         |      START TIME      |       END TIME       | USER NAME |
+--------------------------------------+--------------------------------
+------------------------+----------------------+----------------------+-----------+
| b3f09f23-5814-4ac1-95d0-45b1c1e36ac6 | fabric configure:Validate
| Completed(18.645056ms)  | 2024-10-30T09:37:15Z | 2024-10-30T09:37:15Z | user      |
|                                      | Fabric
|                        |                      |                      |           |
| dc6acd7e-1eb4-4a72-aed1-24d959d7bfee | fabric configure:Validate
| Completed(11.040851ms)  | 2024-10-30T09:36:56Z | 2024-10-30T09:36:56Z | user      |
|                                      | Fabric
|                        |                      |                      |           |
| d84923a1-d995-46de-9a9c-6a297fe32fcb | Debug fabric
| Completed(215.010696ms) | 2024-10-30T09:32:57Z | 2024-10-30T09:32:57Z | user      |
|                                      | health:DebugFabricHealth
|                        |                      |                      |           |
| 5465a2ad-d880-41f2-871e-d9225ccb7135 | Debug enable-smartdrc
| Completed(2.774389ms)   | 2024-10-30T09:31:14Z | 2024-10-30T09:31:14Z | user      |
```

```
|                                    | config:EnableSmartDrcConfig
|                          |                    |                      |           |
| c4233e76-1123-48a9-95ed-38b71cec70a8 | fabric device show
| Failed(6.763279ms)       | 2024-10p-30T09:31:08Z | 2024-10-30T09:31:08Z |           |
| ea4ac2a7-297a-4ae7-a1d6-a10bdcfe32f3 | Debug enable-smartdrc
| Completed(3.284577ms)    | 2024-10-30T09:31:01Z | 2024-10-30T09:31:01Z | user      |
|                                    | config:EnableSmartDrcConfig
|                          |                    |                      |           |
| 3c6f093c-026f-4cac-a39d-3efb515c88c0 | intial device discovery event
| Completed(7.387153ms)    | 2024-10-30T09:30:56Z | 2024-10-30T09:30:56Z |           |
| 0b5278f2-5b71-4e46-93d6-6d675c2ab42f | config show
| Completed(12.34922118s) | 2024-10-30T09:28:32Z | 2024-10-30T09:28:44Z | user      |
| a8ce8e93-aad1-46a0-af94-a2375349dc39 | fabric show --name
| Completed(65.45142ms)    | 2024-10-30T09:26:08Z | 2024-10-30T09:26:08Z | user      |
+--------------------------------------+---------------------------------
+------------------------+---------------------+----------------------+-----------+
```

- Display the list of event histories.

```
efa fabric execution show-event
+--------------------------------------+----------------------
+---------+-----------------+--------+--------------+----------------
+---------------------------------+-------+
|            EXECUTION UUID            |
DATE          | SERVICE |     EVENT
| DEVICE | MESSAGE-TYPE | MESSAGE-OBJECT |            MESSAGE              | ERROR |
+--------------------------------------+----------------------
+---------+---------------+--------+--------------+----------------
+---------------------------------+-------+
| ed5407c3-a42d-46ed-8cf2-829801aab90a | 2024-10-29T19:00:26Z
| fabric  | fabric-create |
| fabric       | default        | Validation succeeded for       |       |
|                                       |                       |
|               |        |              |
| Fabric default                 |       |
| ed5407c3-a42d-46ed-8cf2-829801aab90a | 2024-10-29T19:00:26Z
| fabric  | fabric-create |
| fabric       | default        | Allocation succeeded for       |       |
|                                       |                       |
|               |        |              |
| Fabric default                 |       |
| ed5407c3-a42d-46ed-8cf2-829801aab90a | 2024-10-29T19:00:26Z
| fabric  | fabric-create |
| fabric       | default        | Configuration on devices       |       |
|                                       |                       |
|               |        |              |
| succeeded for Fabric default   |       |
| ef02fc6f-9fe0-4587-885a-592cc34887c5 | 2024-10-29T19:02:32Z
| fabric  | fabric-create |
| fabric       | fs             | Validation succeeded for       |       |
|                                       |                       |
|               |        |              |
| Fabric fs                      |       |
| ef02fc6f-9fe0-4587-885a-592cc34887c5 | 2024-10-29T19:02:33Z
| fabric  | fabric-create |
| fabric       | fs             | Allocation succeeded for       |       |
|                                       |                       |
|               |        |              |
| Fabric fs                      |       |
| ef02fc6f-9fe0-4587-885a-592cc34887c5 | 2024-10-29T19:02:33Z
| fabric  | fabric-create |
| fabric       | fs             | Configuration on devices       |       |
|                                       |                       |
|               |        |              |
| succeeded for Fabric fs        |       |
```

```
+-------------------------------------+---------------------
+---------+---------------+--------+--------------+----------------
+---------------------------------+-------+
--- Time Elapsed: 22.160744ms ---
```

- Delete execution entries older than the specified days.

```
efa fabric execution delete --days 0
Deletion Succeeded
--- Time Elapsed: 27.156452ms ---
```

11. To display the IP fabric configuration, run the fabric show config command.

```
efa fabric show-config --name fs --ip 10.64.188.59
{{
  "fabric_name": "fs",
  "switches": [
    {
      "switch": "10.64.188.59",
      "role": "Leaf",
      "GlobalProperties": {
        "OverlayGwBroadcastLocalBias": false
      },
      "bgp": {
        "local_as": "4200000100",
        "network": "172.31.66.35/32",
        "max_paths": "0",
        "ibg_max_paths": "6",
        "ebgp_max_paths": "6",
        "bfd_rx": "0",
        "bfd_tx": "0",
        "bfd_multiplier": "0",
        "peer_groups": [
          {
            "name": "fsmlag_peergroup_ibgp",
            "description": "peer containing mlag bgp peer",
            "bfd": "true",
            "remote_as": "4200000100",
            "password": "fs",
            "update_source": "",
            "next_hop_self": true,
            "fast_ext_failover": true,
            "bfd_profile": "fs"
          }
        ],
        "neighbors": [
          {
            "remote_ip": "10.15.16.2",
            "remote_as": "4200000100",
            "peer_group": "fsmlag_peergroup_ibgp",
            "bgp_multihop": "",
            "enable_peer_as_check": false,
            "bfd_enable": false,
            "status": "ESTAB",
            "is_activated": true,
            "next_hop_self": true,
            "up_time": "44s",
            "neighbor_afi": "ipv4",
            "neighbor_safi": "unicast",
            "password": ""
          }
        ],
        "l2vpn": {
          "graceful_restart": true,
          "evpn_neighbors": []
        },
```

```
        "afnetwork": null,
        "router_id": "172.32.128.99",
        "listen_limit": 0
      },
      "router_id": "172.32.128.99",
      "overlay_gateway": null,
      "overlay_network": {
        "nvo": [
          {
            "name": "base",
            "member_nve": [
              "base"
            ],
            "member_vni_domain": "base",
            "split_horizon": "base"
          },
          {
            "name": "fs",
            "member_nve": [
              "fs"
            ],
            "member_vni_domain": "fs",
            "split_horizon": "fs"
          }
        ],
        "nve": [
          {
            "name": "base",
            "source_interface ": ""
          },
          {
            "name": "fs",
            "source_interface ": "loopback 2"
          }
        ],
        "vni_domain": [
          {
            "name": "base",
            "auto_offset": 0
          },
          {
            "name": "fs",
            "auto_offset": 0
          }
        ]
      },
      "evpn": {
        "name": "fs",
        "duplicate_mac_timer_value": "10",
        "max_count": "8",
        "rd": "",
        "overlay_nve": "fs",
        "dampening_time": "6",
        "route_target": [
          {
            "ignore_as": true
          }
        ]
      },
      "router_pim": [],
      "ip_prefix_list": [],
      "interfaces": [
        {
          "interface_name": "0/1",
```

```
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "clusterPeerIntfMember",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "0/2",
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "clusterPeerIntfMember",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "0/3",
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "Ethernet 0/3",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "0/4",
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "Ethernet 0/4",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "0/5",
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "Ethernet 0/5",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "0/6",
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "Ethernet 0/6",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "0/7",
          "interface_type": "ethernet",
          "interface_ip_address": "",
          "interface_description": "Ethernet 0/7",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "62",
          "interface_type": "port-channel",
          "interface_ip_address": "10.15.16.3/31",
          "interface_description": "MCTPeerInterface",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "1",
          "interface_type": "loopback",
          "interface_ip_address": "172.32.128.99/32",
          "interface_description": "",
          "ip_pim_sparse": false
        },
        {
          "interface_name": "2",
          "interface_type": "loopback",
```

```
          "interface_ip_address": "172.31.66.35/32",
          "interface_description": "",
          "ip_pim_sparse": false
        }
      ],
      "vrfs": [
        {
          "vrf": "default-vrf",
          "vrf_interfaces": [
            {
              "interface_type": "loopback",
              "interface_name": "1"
            },
            {
              "interface_type": "loopback",
              "interface_name": "2"
            },
            {
              "interface_type": "port-channel",
              "interface_name": "62"
            }
          ]
        }
      ]
    },
      ]
}}
--- Time Elapsed: 12.378764264s ---
```

12. Validate the physical topology of an IP fabric.

```
efa fabric topology validate physical --name fs
Fabric Name: fs, Fabric Type: non-clos
Fabric fs is healthy
--- Time Elapsed: 29.856064ms ---
```

13. Create a clone of the existing fabric.

```
(efa:extreme)extreme@node-1:~$ efa fabric clone --source fabric1 --destination
BLR_FABRIC
Clone Fabric [Success]

--- Time Elapsed: 2.889625146s ---
(efa:extreme)extreme@node-1:~$ efa fabric show --name fabric1

Fabric Name: fabric1, Fabric Description: , Fabric Type: non-clos, Fabric Status:
configure-success, Fabric Health: Green
+--------------+-------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
|  IP ADDRESS  | RACK  | HOST NAME |    ASN     | ROLE | DEVICE STATE
|  APP STATE  | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+--------------+-------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+
| 10.64.224.17 | Rack1 | Leaf-17   | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 12      | 11    |
| 10.64.224.16 | Rack1 | Leaf-16   | 4200000000 | leaf | provisioned
| cfg in-sync | NA                | NA              | 12      | 11    |
+--------------+-------+-----------+------------+------+--------------+-------------
+-------------------+-----------------+---------+-------+

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, BFDU/BFDD
- BFD Update/Delete, VRFU - VRF Update
DD - Dependent Device Update, DA/DR - Device Add/ReAdd, ASN - Asn
```

```
Update, SYS - System Properties Update
MD5 - BGP MD5 Password, BGPU/BGPD - Router BGP Update/Delete, BGPLL
- BGP Listen Limit, POU - Port Channel Update, BGPAFD/BGPAFU/BGPAFC - BGP Address
family Delete/Update/Create, L2VPNC/L2VPNU/L2VPND - BGP L2VPN Address family Create/
Update/Delete , BGPPGU/BGPPGD - BGP Peer Group Update/Delete, NA - Not Applicable

PENDING CONFIGS:
MCT - MCT Cluster, O - Overlay Gateway, SYSP - System Properties,
INTIP - Interface IP, BGP - Router BGP, VRFM - VRF Member, BGPPG - BGP Peer Group
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: 71.517899ms ---


(efa:extreme)extreme@node-1:~$ efa fabric show --name BLR_FABRIC

Fabric Name: BLR_FABRIC, Fabric Description: , Fabric Type: non-clos, Fabric
Status: created, Fabric Health: Green
+------------+------+-----------+-----+------+--------------+-----------
+--------------------+-----------------+---------+-------+
| IP ADDRESS | RACK | HOST NAME | ASN | ROLE
| DEVICE STATE | APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+------------+------+-----------+-----+------+--------------+-----------
+--------------------+-----------------+---------+-------+
+------------+------+-----------+-----+------+--------------+-----------
+--------------------+-----------------+---------+-------+
--- Time Elapsed: 62.731885ms ---
```

14. Run the **efa fabric topology show underlay --name <fabric-name>** command to get the latest status of the BGP session.
    a. Verify the cluster configuration by XCO.

```
mlag
  bringup-delay 450
  peer efa_fabric-cluster-1
    keepalive efa_fabric-cluster-1
      type primary
      destination 10.20.20.4
      source-interface port-channel 64
    keepalive efa_fabric-cluster-1-sec
      type secondary
    !
```

```
  !
!
```

SLX Configuration	OS ONE Configuration

cluster efa_fabric-cluster-1
 peer 10.20.20.2
 peer-interface Port-channel 64
 peer-keepalive
  auto
 !
 member vlan all
 member bridge-domain all
!
	mlag
	  bringup-delay 450
	  peer efa_fabric-cluster-1
	  keepalive efa_fabric-cluster-1
	  type primary
	  destination 10.20.20.7
	  source-interface port-channel 64
	  keepalive efa_fabric-cluster-1-sec
	  type secondary
	  !
	  !
	For OS ONE, configure the cluster under MLAG. SLX cluster name is equivalent to the peer name in OS ONE.
	Peer IP and source interface is configured under keepalive for OS ONE.

b. Verify the BGP configurations on SLX and OS ONE devices.:

SLX configurations	OS ONE configurations

router bgp
local-as 4200000001
	router bgp
	  local-as 4200000000
	  as-notation as-plain

capability as4-enable listen-limit 101	capability as4-enable, not available listen-limit, not supported
fast-external-fallover	fast-external-fallover, under peer-group level

neighbor 10.20.20.2 remote-as 4200000001
 neighbor 10.20.20.2 next-hop-self
 neighbor 10.20.20.2 password 9YJtsSpf==
 neighbor 10.20.20.2 bfd
	peer-group FAB_mlag_peergroup_ibgp
	  description peer containing mlag bgp peer
	  remote-as 4200000000
	  fast-external-failover
	  address-family ipv4 unicast
	  nexthop-self
	  activate
	  !
	  neighbor 10.20.20.2
	  !
	To configure the OS ONE neighbors, create a peer-group and assign the neighbor properties under peer group, And define the neighbor under peer-group.

SLX configurations	OS ONE configurations

address-family ipv4 unicast
  maximum-paths 9
  graceful-restart
 !
	aaddress-family ipv4 unicast
	  graceful-restart
	  use-multiple-paths ibgp
	maximum-paths 9
	  use-multiple-paths ebgp
	maximum-paths 9
	  activate
	 !

address-family l2vpn evpn
  graceful-restart
 !
	address-family l2vpn evpn
	  graceful-restart
	  activate
	 !

Configure Tenant Services

You can configure tenant services on the Extreme OS ONE devices.

About This Task

Follow this procedure to complete the tenant provisioning on day-1 to support OS ONE devices.

Note

For information on command syntax and parameter description, see *ExtremeCloud Orchestrator Command Reference Guide*.

Procedure

1. Configure tenant LCM.

   The changes are highlighted that support enabling OS ONE device.

   a. Create a tenant.

```
efa tenant create

      --name string                          Name of the Tenant
      --description string                   Description of the Tenant
      --type string                          Default tenant type is set to
private. Valid values are shared | private. POs of shared tenant are available to
other tenants. Cannot create VRF, EndpointGroup, BGP Peer and BGP Peer-Group for
the shared tenants
      --l2-vni-range string                  [applicable only for Non auto-VNI
fabric] Contiguous Range of L2 VNIs in ascending order will be reserved for the
tenant within the scope of a fabric.
                                             Valid VNI values are <1-16777215>.
Example: 14201-14300
      --l3-vni-range string                  [applicable only for Non auto-VNI
fabric] Contiguous Range of L3 VNIs in ascending order will be reserved for the
tenant within the scope of a fabric.
                                             Valid VNI values are <1-16777215>.
Example: 14201-14300
      --vlan-range string                    Contiguous Range of VLANs to be
reserved for the tenant. Valid values are <2-4090>. Example: 2-100
      --vrf-count int                        Number of VRFs reserved for the Tenant
      --enable-bd                            Enable BD capability for networks
created under this Tenant. Example: --enable-bd=true | --enable-bd=false
      --port stringArray                     List of physical ports
of devices which will be reserved as data ports. Example:
```

```
SW1_IP[0/1,0/2,0/12-15,0/5:4],SW2_IP[0/1,0/4,0/5:1-2,0/5:3,0/9-20]
      --mirror-destination-port stringArray   [traffic-mirroring is not supported]
List of physical ports of devices which will be reserved as mirror destination
ports.
                                              Example:
SW1_IP[0/1,0/2,0/12-15,0/5:4],SW2_IP[0/1,0/4,0/5:1-2,0/5:3,0/9-20]
```

b. Update a tenant.

```
efa tenant update

--name string                            Name of the Tenant
      --operation string                      Operation code. Valid values are desc-
update | vni-update | port-add | port-delete | vlan-add | vlan-delete | vlan-update
| num-vrf-update |
                                               enable-bd-update | mirror-
destination-port-add | mirror-destination-port-delete
      --description string                    Description of the Tenant
      --l2-vni-range string                   [Applicable for non-auto vni fabric
only] Contiguous Range of L2 VNIs in ascending order will be reserved for the
tenant within the scope of a fabric.
                                              Valid VNI values are <1-16777215>.
Example: 14201-14300
      --l3-vni-range string                   [Applicable for non-auto vni fabric
only]Contiguous Range of L3 VNIs in ascending order will be reserved for the tenant
within the scope of a fabric.
                                              Valid VNI values are <1-16777215>.
Example: 14201-14300
      --vlan-range string                     Contiguous Range of VLANs to be
reserved for the tenant. Valid values are <2-4090>. Example: 2-100
      --vrf-count int                         Number of VRFs reserved for the Tenant
      --enable-bd                             BD capability for networks created
under this Tenant. Example --enable-bd=true | --enable-bd=false
      --port stringArray                      List of physical ports
of devices which will be reserved as data ports. Example:
SW1_IP[0/1,0/2,0/12-15,0/5:4],SW2_IP[0/1,0/4,0/5:1-2,0/5:3,0/9-20]
      --mirror-destination-port stringArray   List of physical ports of devices
which will be reserved as mirror destination ports.
                                              Example:
SW1_IP[0/1,0/2,0/12-15,0/5:4],SW2_IP[0/1,0/4,0/5:1-2,0/5:3,0/9-20]
      --force                                 Force option which is valid only for
port-delete operation is deprecated.
```

c. Delete a tenant.

There are no changes to the **`efa tenant delete`** and **`efa tenant show`** command.

```
efa tenant delete

--name string   Name of the Tenant
      --force         Force tenant deletion if the option is provided
```

2. Configure EPG LCM.

EPG creation is supported for L2 EPGs on both BD/VLAN tenant network and Auto-VNI fabric. The changes are highlighted that support enabling OS ONE device.

a. Create an EPG.

The highlighted Layer 2 options are not supported for OS ONE devices, and appropriate errors are displayed for any attempt to configure them via the CLI or REST. The highlighted layer-3 and VNI options are not supported.

```
efa tenant epg create

--name string                                    Name of the EndpointGroup
      --tenant string                                 Name of the Tenant
      --description string                            Description of the
```

```
EndpointGroup
      --port stringArray                                Device IP along
with ethernet port details. Example: SW1_IP[0/1], SW2_IP[0/5,0/6],
SW3_IP[0/7:1-2,0/7:4,0/10]
      --po stringArray                                  List of PortChannels.
Example: po1,po2
      --switchport-mode string                          Configures switchport mode
on the interfaces. Valid values are access |trunk | trunk-no-default-native.
Default value is set to trunk
      --type string                                     Configures Bgp service
type. Valid values are l3-hand-off | extension | port-profile. Default value is set
to extension
      --switchport-native-vlan-tagging                  Enable the native vlan
characteristics on the ports of this endpoint group. Valid only if mode is set to
trunk
      --single-homed-bfd-session-type string            Sets BFD session type on
ports of this endpoint group. Valid values are auto | hardware | software
      --switchport-native-vlan string                   Configures native Vlan on
the interfaces. Valid values are 2-4090
      --ctag-range string                               Customer Vlan range in
comma and hyphen separated format. Example: 2-20,30,40,50-55
      --vrf string                                      Vrf to which these networks
are attached. Example: VRF-GREEN
      --l3-vni string                                   L3 VNI to be used for this
Vrf. Valid VNI values are <1-16777215>. Refer to associated tenant for valid VNI
range. Example: 1002
      --l2-vni stringArray                              L2 VNI to be used for this
network in the format ctag:l2-vni. Refer to associated tenant for valid VNI range.
Example: 10:12156
      --ip-mtu stringArray                              IP MTU to be used for this
network in the format ctag:ip-mtu. Valid values are <1280-9194>. Example: 10:9190
      --anycast-ip stringArray                          IPv4 anycast address in the
format ctag:anycast-ip. Example 10:10.1.1.1/24
      --anycast-ipv6 stringArray                        IPv6 anycast address in the
format ctag:anycast-ipv6. Example: 10:2000::1/64
      --local-ip stringArray                            IPv4 local address in the
format ctag,device-ip:local-ip
      --local-ipv6 stringArray                          IPv6 local address in the
format ctag,device-ip:local-ipv6
      --ip-icmp-redirect stringArray                    Sets IPv4 icmp redirect
flag in the format ctag:icmp-redirect. Example: 1002:true
      --ipv6-icmp-redirect stringArray                  Sets IPv6 icmpv6 redirect
flag in the format ctag:icmpv6-redirect. Example: 1002:true
      --dhcpv4-relay-address-ip stringArray             IPv4 dhcp relay address in
the format ctag,device-ip:dhcp-address-ip. Example 10,10.20.20.1:10.1.1.1
      --dhcpv6-relay-address-ip stringArray             IPv6 dhcp relay address in
the format ctag,device-ip:dhcp-address-ipv6. Example 10,10.20.20.1:3::3
      --dhcpv4-relay-address-ip-vrf stringArray         IPv4 dhcp relay address
ip vrf in the format ctag,device-ip:dhcp-address-ip,vrf-name. Example
10,10.20.20.1:10.1.1.1,vrf1
      --dhcpv6-relay-address-ip-vrf stringArray         IPv6 dhcp relay address
ip vrf in the format ctag,device-ip:dhcp-address-ipv6,vrf-name. Example
10,10.20.20.1:3::3,vrf1
      --dhcpv6-relay-address-ip-interface stringArray   IPv6 dhcp relay
address ip interface in the format ctag,device-ip:dhcp-address-
ipv6,interfaceType,interfaceName.
                                                        Example
10,10.20.20.1:3::3,eth,0/21 | 10,10.20.20.1:3::3,po,po1
      --dhcpv4-relay-gateway-ip stringArray             IPv4 dhcp relay gateway in
the format ctag,device-ip:dhcp-gateway-ip. Example 10,10.20.20.1:10.1.1.1
      --dhcpv4-relay-gateway-ip-interface stringArray   IPv4 dhcp relay gateway ip
interface in the format ctag,device-ip:dhcp-gateway-ip,interfaceType,interfaceName.
                                                        Example
10,10.20.20.1:10.1.1.1,eth,0/21
```

```
      --dhcpv4-relay-gateway-interface stringArray      IPv4 dhcp relay gateway
interface in the format ctag,device-ip:interfaceType,interfaceName.
                                                        Example
10,10.20.20.1:eth,0/21
      --dhcpv6-relay-gateway-interface stringArray      IPv6 dhcp relay gateway
interface in the format ctag,device-ip:interfaceType,interfaceName.
                                                        Example
10,10.20.20.1:eth,0/21
      --dhcpv6-relay-gateway-interface-ip stringArray   IPv6 dhcp relay gateway
interface ipv6 in the format ctag,device-ip:interfaceType,interfaceName,dhcp-
gateway-ipv6.
                                                        Example
10,10.20.20.1:eth,0/21,3::3
      --bridge-domain stringArray                       Name of the bridge-domain
in the format ctag:bridge-domain. Example: 1002:Ten1-Net2
      --ipv6-nd-mtu stringArray                         Sets Mtu of IPv6 ND Valid
values are <1280-65535>. Format ctag:mtu. Example: 1002:1800
      --ipv6-nd-managed-config stringArray              Sets managed config flag in
IPv6 Router advertisement. Format ctag:managedflag. Example: 1002:true
      --ipv6-nd-other-config stringArray                Sets other config flag in
IPv6 Router advertisement. Format ctag:otherflag. Example: 1002:true
      --ipv6-nd-prefix stringArray                      Configures IPv6 prefix
address. Format ctag:prefix1,prefix2,. Example: 2000:1:5::/64,1:6::/64
      --ipv6-nd-prefix-valid-lifetime stringArray       Sets IPv6 prefix valid
lifetime. Valid range <0-4294967295>/infinite. Format ctag,prefix:validTime
                                                        Example: 2000,1:5::/64:1500
or 2000,5::/64:infinite
      --ipv6-nd-prefix-preferred-lifetime stringArray   Sets IPv6 prefix preferred
lifetime. Valid range <0-4294967295>/infinite. Format ctag,prefix:preferredTime
                                                        Example: 2000,1:5::/64:1500
or 2000,1:5::/64:infinite
      --ipv6-nd-prefix-no-advertise stringArray         Enable to prevent prefix
advertisement. Format ctag,prefix:noadvertiseflag. Example: 1002,1:5::/64:true
      --ipv6-nd-prefix-config-type stringArray          Set config type for
IPv6 prefix. Valid values are no-autoconfig| no-onlink | off-link. Format
ctag,prefix:configType
                                                        Example: 1002,1:5::/64:no-
autoconfig/1002,1:5::/64:no-onlink/1002,1:5::/64:offlink
      --ctag-description stringArray                    Description of vlan in the
format ctag:Description. Example: 1002:VlanDescription
      --suppress-arp stringArray                        Sets suppress-arp flag
to this network. Format ctag:suppress-arp. Example: --suppress-arp 1002:true  --
suppress-arp 1003:false
      --suppress-nd stringArray                         Sets suppress-nd flag to
this network. Format ctag:suppress-nd. Example: --suppress-nd 1002:true  --suppress-
nd 1003:false
      --pp-mac-acl-in stringArray                       Apply MAC ACL in ingress
direction on ethernet / portchannel interfaces. The only supported ACL name is ext-
mac-permit-any-mirror-acl
                                                        and only supported ACL type
is extended. Format --pp-mac-acl-in <acl-name> Example: --pp-mac-acl-in ext-mac-
permit-any-mirror-acl
      --pp-mac-acl-out stringArray                      Apply MAC ACL in egress
direction on ethernet / portchannel interfaces. The only supported ACL name is ext-
mac-permit-any-mirror-acl and
                                                        only supported ACL type
is extended. Format --pp-mac-acl-out <acl-name> Example: --pp-mac-acl-out ext-mac-
permit-any-mirror-acl
      --pp-ip-acl-in stringArray                        Apply IP ACL in ingress
direction on ethernet / portchannel interfaces. The only supported ACL name is ext-
ip-permit-any-mirror-acl
                                                        and only supported ACL type
is extended. Format --pp-ip-acl-in <acl-name> Example: --pp-ip-acl-in ext-ip-permit-
any-mirror-acl
```

```
      --pp-ip-acl-out stringArray                          Apply IP ACL in egress
direction on ethernet / portchannel interfaces. The only supported ACL name is ext-
ip-permit-any-mirror-acl
                                                           and only supported ACL type
is extended. Format --pp-ip-acl-out <acl-name> Example: --pp-ip-acl-out ext-ip-
permit-any-mirror-acl
      --pp-ipv6-acl-in stringArray                         Apply IPv6 ACL in ingress
direction on ethernet / portchannel interfaces. The only supported ACL name is ext-
ipv6-permit-any-mirror-acl
                                                           and only supported ACL type
is extended. Format --pp-ipv6-acl-in <acl-name> Example: --pp-ipv6-acl-in ext-ipv6-
permit-any-mirror-acl
      --np-mac-acl-in stringArray                          Apply MAC ACL in ingress
direction on vlan. The only supported ACL name is ext-mac-permit-any-mirror-acl
                                                           and only supported ACL type
is extended. Format --np-mac-acl-in <ctag:acl-name> Example: --np-mac-acl-in
101:ext-mac-permit-any-mirror-acl
      --np-mac-acl-out stringArray                         Apply MAC ACL in egress
direction on vlan. The only supported ACL name is ext-mac-permit-any-mirror-acl
                                                           and only supported ACL type
is extended. Format --np-mac-acl-out <ctag:acl-name> Example: --np-mac-acl-out
101:ext-mac-permit-any-mirror-acl
      --np-ip-acl-in stringArray                           Apply IP ACL in ingress
direction on ve interface. The only supported ACL name is ext-ip-permit-any-mirror-
acl
                                                           and only supported ACL type
is extended. Format --np-ip-acl-in <ctag:acl-name> Example: --np-ip-acl-in 101:ext-
ip-permit-any-mirror-acl
      --np-ip-acl-out stringArray                          Apply IP ACL in egress
direction on ve interface. The only supported ACL name is ext-ip-permit-any-mirror-
acl
                                                           and only supported ACL type
is extended. Format --np-ip-acl-out <ctag:acl-name> Example: --np-ip-acl-out
101:ext-ip-permit-any-mirror-acl
      --np-ipv6-acl-in stringArray                         Apply IPv6 ACL in ingress
direction on ve interface. The only supported ACL name is ext-ipv6-permit-any-
mirror-acl
                                                           and only supported ACL type
is extended. Format --np-ipv6-acl-in <ctag:acl-name> Example: --np-ipv6-acl-in
101:ext-ipv6-permit-any-mirror-acl
```

b. Update an EPG.

```
efa tenant epg update

--name string                                      Name of the EndpointGroup
      --tenant string                              Name of the Tenant
      --operation string                           Operation code. Valid
values are port-group-add | port-group-delete | port-property-add | port-property-
delete | port-property-update |
                                                   ctag-range-add | ctag-range-
delete | network-property-add | network-property-delete | network-property-update |
                                                   vrf-add | vrf-delete |
local-ip-add | local-ip-delete | anycast-ip-add | anycast-ip-delete
                                                   dhcp-relay-address-ip-add |
dhcp-relay-address-ip-delete | dhcp-relay-gateway-ip-add | dhcp-relay-gateway-ip-
delete
      --port stringArray                           List of physical ports of
device on which tenant network will be configured
                                                   Example:
SW1_IP[0/1,0/2,0/12-15,0/5:4],SW2_IP[0/1,0/4,0/5:1-2,0/5:3,0/9-20]
      --po stringArray                             List of PortChannels on
which tenant network will be configured. Example: po1 or po1,po2
      --switchport-mode string                     Configures switchport mode
on the interfaces. Valid values are access |trunk | trunk-no-default-native.
```

```
Default value is set to trunk
      --switchport-native-vlan string                    Configures native vlan on
the interfaces. Valid values are 2-4090
      --switchport-native-vlan-tagging                   Enable native-vlan
characteristics on the ports of this endpoint group. Valid only if mode is set to
trunk
      --single-homed-bfd-session-type string             Sets BFD session type on
ports of this endpoint group. Valid values are auto | hardware | software
      --pp-mac-acl-in stringArray                        Apply MAC ACL for mirror
action in ingress direction on ethernet / portchannel interfaces. The only
supported ACL name is
                                                         ext-mac-permit-any-mirror-
acl and only supported ACL type is extended. Format --pp-mac-acl-in <acl-name>
                                                         Example: --pp-mac-acl-in
ext-mac-permit-any-mirror-acl
      --pp-mac-acl-out stringArray                       Apply MAC ACL for mirror
action in egress direction on ethernet / portchannel interfaces. The only supported
ACL name is
                                                         ext-mac-permit-any-mirror-
acl and only supported ACL type is extended. Format --pp-mac-acl-out <acl-name>
                                                         Example: --pp-mac-acl-out
ext-mac-permit-any-mirror-acl
      --pp-ip-acl-in stringArray                         Apply IP ACL for mirror
action in ingress direction on ethernet / portchannel interfaces. The only
supported ACL name is
                                                         ext-ip-permit-any-mirror-
acl and only supported ACL type is extended. Format --pp-ip-acl-in <acl-name>
                                                         Example: --pp-ip-acl-in ext-
ip-permit-any-mirror-acl
      --pp-ip-acl-out stringArray                        Apply IP ACL for mirror
action in egress direction on ethernet / portchannel interfaces. The only supported
ACL name is
                                                         ext-ip-permit-any-mirror-
acl and only supported ACL type is extended. Format --pp-ip-acl-out <acl-name>
                                                         Example: --pp-ip-acl-out
ext-ip-permit-any-mirror-acl
      --pp-ipv6-acl-in stringArray                       Apply IPv6 ACL for mirror
action in ingress direction on ethernet / portchannel interfaces. The only
supported ACL name is
                                                         ext-ipv6-permit-any-mirror-
acl and only supported ACL type is extended. Format --pp-ipv6-acl-in <acl-name>
                                                         Example: --pp-ipv6-acl-in
ext-ipv6-permit-any-mirror-acl
      --ctag-range string                                Customer vlan range in
comma and hyphen separated format. Example: 2-20,30,40,50-55
      --vrf string                                       Vrf to which these networks
are attached. Example: VRF-GREEN
      --l3-vni string                                    L3 VNI to be used for this
Vrf
      --l2-vni stringArray                               L2 VNI to be used for this
network in the format ctag:l2-vni. Example: 10:12156
      --ip-mtu stringArray                               IP MTU to be used for this
network in the format ctag:ip-mtu. Valid values are <1280-9194>. Example: 10:9190
      --anycast-ip stringArray                           IPv4 anycast address in the
format ctag:anycast-ip. Example 10:10.1.1.1/24
      --anycast-ipv6 stringArray                         IPv6 anycast address in the
format ctag:anycast-ipv6. Example: 10:2000::1/64
      --bridge-domain stringArray                        Bridge domain name in the
format ctag:bridge-domain. Example: 1002:Ten1-Net2
      --ctag-description stringArray                     Vlan/BD description in the
format ctag:vlandescription. Example: 1002:VlanDescription
      --local-ip stringArray                             IPv4 local address in the
format ctag,device-ip:local-ip
      --local-ipv6 stringArray                           IPv6 local address in the
```

```
format ctag,device-ip:local-ipv6
      --ip-icmp-redirect stringArray                    Sets IPv4 icmp redirect
flag in the format ctag:icmp-redirect. Example: --ip-icmp-redirect 1002:true --ip-
icmp-redirect 1003:false
      --ipv6-icmp-redirect stringArray                  Sets IPv6 icmpv6 redirect
flag in the format ctag:icmpv6-redirect. Example: --ipv6-icmp-redirect 1002:true --
ipv6-icmp-redirect 1003:false
      --ipv6-nd-mtu stringArray                         Sets Mtu of IPv6 ND. Valid
values are <1280-65535>. Format ctag:mtu
      --ipv6-nd-managed-config stringArray              Managed config flag to be
set in Router advertisement. Format ctag:ipv6-nd-managed-config
      --ipv6-nd-other-config stringArray                Other config flag to be set
in Router advertisement. Format ctag:ipv6-nd-other-config
      --ipv6-nd-prefix stringArray                      Configures IPv6 prefix
address. Format ctag:prefix1,prefix2,. Example: 2000:1:5::/64,1:6::/64
      --ipv6-nd-prefix-valid-lifetime stringArray       Sets IPv6 prefix valid
lifetime. Valid range <0-4294967295>/infinite. Format ctag,prefix:validTime
                                                        Example: 2000,1:5::/64:1500
or 2000,5::/64:infinite
      --ipv6-nd-prefix-preferred-lifetime stringArray   Sets IPv6 prefix preferred
lifetime. Valid range <0-4294967295>/infinite. Format ctag,prefix:preferredTime
                                                        Example: 2000,1:5::/64:1500
or 2000,1:5::/64:infinite
      --ipv6-nd-prefix-no-advertise stringArray         Enable to prevent prefix
advertisement. Format ctag,prefix:noadvertiseflag. Example: 1002,1:5::/64:true
      --ipv6-nd-prefix-config-type stringArray          Set config type for
IPv6 prefix. Valid values are no-autoconfig| no-onlink | off-link. Format
ctag,prefix:configType
                                                        Example: 1002,1:5::/64:no-
autoconfig/1002,1:5::/64:no-onlink/1002,1:5::/64:off-link
      --suppress-arp stringArray                        Sets suppress-arp flag
to this network. Format ctag:suppress-arp. Example: --suppress-arp 1002:true --
suppress-arp 1003:false
      --suppress-nd stringArray                         Sets suppress-nd flag to
this network. Format ctag:suppress-nd. Example: --suppress-nd 1002:true --suppress-
nd 1003:false
      --np-mac-acl-in stringArray                       Apply MAC ACL for mirror
action in ingress direction on vlan. The only supported ACL name is ext-mac-permit-
any-mirror-acl
                                                        and only supported ACL type
is extended. Format --np-mac-acl-in <ctag:acl-name>
                                                        Example: --np-mac-acl-in
101:ext-mac-permit-any-mirror-acl
      --np-mac-acl-out stringArray                      Apply MAC ACL for mirror
action in egress direction on vlan. The only supported ACL name is ext-mac-permit-
any-mirror-acl
                                                        and only supported ACL type
is extended. Format --np-mac-acl-out <ctag:acl-name>
                                                        Example: --np-mac-acl-out
101:ext-mac-permit-any-mirror-acl
      --np-ip-acl-in stringArray                        Apply IP ACL for mirror
action in ingress direction on ve interface. The only supported ACL name is ext-ip-
permit-any-mirror-acl
                                                        and only supported ACL type
is extended. Format --np-ip-acl-in <ctag:acl-name>
                                                        Example: --np-ip-acl-in
101:ext-ip-permit-any-mirror-acl
      --np-ip-acl-out stringArray                       Apply IP ACL for mirror
action in egress direction on ve interface. The only supported ACL name is ext-ip-
permit-any-mirror-acl
                                                        and only supported ACL type
is extended. Format --np-ip-acl-out <ctag:acl-name>
                                                        Example: --np-ip-acl-out
101:ext-ip-permit-any-mirror-acl
```

```
      --np-ipv6-acl-in stringArray                        Apply IPv6 ACL for mirror
action in ingress direction on ve interface. The only supported ACL name is ext-
ipv6-permit-any-mirror-acl
                                                          and only supported ACL type
is extended. Format --np-ipv6-acl-in <ctag:acl-name>
                                                          Example: --np-ipv6-acl-in
101:ext-ipv6-permit-any-mirror-acl
      --dhcpv4-relay-address-ip stringArray               IPv4 dhcp relay address in
the format ctag,device-ip:dhcp-address-ip. Example 10,10.20.20.1:10.1.1.1
      --dhcpv6-relay-address-ip stringArray               IPv6 dhcp relay address in
the format ctag,device-ip:dhcp-address-ipv6. Example 10,10.20.20.1:3::3
      --dhcpv4-relay-address-ip-vrf stringArray           IPv4 dhcp relay address
ip vrf in the format ctag,device-ip:dhcp-address-ip,vrf-name. Example
10,10.20.20.1:10.1.1.1,vrf1
      --dhcpv6-relay-address-ip-vrf stringArray           IPv6 dhcp relay address
ip vrf in the format ctag,device-ip:dhcp-address-ipv6,vrf-name. Example
10,10.20.20.1:3::3,vrf1
      --dhcpv6-relay-address-ip-interface stringArray     IPv6 dhcp relay
address ip interface in the format ctag,device-ip:dhcp-address-
ipv6,interfaceType,interfaceName.
                                                          Example
10,10.20.20.1:3::3,eth,0/21  10,10.20.20.1:3::3,po,po1
      --dhcpv4-relay-gateway-ip stringArray               IPv4 dhcp relay gateway in
the format ctag,device-ip:dhcp-gateway-ip. Example 10,10.20.20.1:10.1.1.1
      --dhcpv4-relay-gateway-ip-interface stringArray     IPv4 dhcp relay gateway ip
interface in the format ctag,device-ip:dhcp-gateway-ip,interfaceType,interfaceName.
                                                          Example
10,10.20.20.1:10.1.1.1,eth,0/21
      --dhcpv4-relay-gateway-interface stringArray        IPv4 dhcp relay gateway
interface in the format ctag,device-ip:interfaceType,interfaceName.
                                                          Example
10,10.20.20.1:eth,0/21
      --dhcpv6-relay-gateway-interface stringArray        IPv6 dhcp relay gateway
interface in the format ctag,device-ip:interfaceType,interfaceName.
                                                          Example
10,10.20.20.1:eth,0/21
      --dhcpv6-relay-gateway-interface-ip stringArray     IPv6 dhcp relay gateway
interface ipv6 in the format ctag,device-ip:interfaceType,interfaceName,dhcp-
gateway-ipv6.
                                                          Example
10,10.20.20.1:eth,0/21,3::3
```

Unsupported switchport mode combinations:

L2 EPG update with the following switchport mode combinations is supported on BD/ Vlan Tenant, and auto-vni:

- `--switchport-mode <trunk|access|trunk-no-default-native>`

c. Delete an EPG.

```
efa tenant epg delete

--name stringArray   EndpointGroup to be deleted. Ex: EPG-1
      --tenant string      Name of the Tenant
      --force              Set to true if EndpointGroup delete requires de-
configuring networks
```

d. Show an EPG.
EPG Show Output for L2 EPG

```
efa tenant epg show --tenant t1 --name epg1
+------+--------+-----------+-------+-----+------------+-------------+------------
+-----+-------+-----------------------------------+
| Name | Tenant |   Type    | Ports | PO  | SwitchPort | Native Vlan |
Ctag Range | Vrf | L3Vni |               State               |
```

```
|      |        |           |       |     |    Mode    |   Tagging   |
|     |       |                                        |
+------+--------+-----------+-------+-----+------------+-------------+------------
+-----+-------+-------------------------------------+
| epg1 |   t1   | extension |       | po2 |   trunk    |    false    |
170     |     |       | epg-with-port-group-and-ctag-range |
|      |        |           |       |     |            |             |
|     |       |                                        |
+------+--------+-----------+-------+-----+------------+-------------+------------
+-----+-------+-------------------------------------+
EndpointGroup Details
PO/Vrf [Flags : * - Unstable]
For 'unstable' entities, run 'efa tenant po/vrf show' for details
--- Time Elapsed: 71.763642ms ---
```

EPG Show Output for L2 EPG in detail

```
efa tenant epg show --tenant t1 --name epg1 --detail
================================================================================================
================================================================================
Name          : epg1
Tenant        : t1
Type          : extension
State         : epg-with-port-group-and-ctag-range
Description   :

Ports         :
POs           : po2
Port Property : SwitchPort Mode               : trunk
              : Native Vlan Tagging           : false
              : Single-Homed BFD Session Type : auto
NW Policy     : Ctag Range                    : 170

+------------+-------------+-----------+------------+-------------+
| MAC ACL IN | MAC ACL OUT | IP ACL IN | IP ACL OUT | IPv6 ACL IN |
+------------+-------------+-----------+------------+-------------+
Port Property ACLs

+------+-------------+-------------+
| Port |  Dev State  |  App State  |
+------+-------------+-------------+
| po2  | 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 |             |             |
+------+----------------------+-------+-------------+-------------+--------------
+----------+----------------------+---------------+--------+---------
+----------------+-------------+-------------+-------------+
| 170  | Tenant L2 Extended BD | 4097  | Auto-BD-4097 |             |
|   F/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: 79.731045ms ---
```

For a brief or detailed output of the VRF for all tenants, a selected tenant, or a selected VRF, see *ExtremeCloud Orchestrator Command Reference Guide*.

3. Configure VRF LCM.

   See VRF Static Route BFD REST Model on page 1190.

Standardization of Northbound API Changes

The OS ONE integration with XCO standardizes Northbound API changes for BGP Peer, BGP Peer Group, and VRF Static Route BFD configurations.

VRF Static Route BFD

Use this topic to learn about configuring BFD for the VRF static route.

Configure VRF Static Route BFD

You can configure BFD for the VRF static route.

About This Task
Follow the procedure to configure BFD for the VRF static route.

Procedure

Run the **`efa tenant vrf create`** command to configure VRF static route BFD.

```
IPv4 Static Route BFD Configuration:
    - The --ipv4-static-route-bfd option is used to configure IPv4 static route BFD.
    - The format for SLXOS is: device-ip,dest-ipv4-addr,source-ipv4-addr[interval,min-
rx,multiplier].
    - Example: --ipv4-static-route-bfd 10.25.25.100,1.1.1.1,2.2.2.2,123,456,3
    - The format for ONEOS is: device-ip,dest-ipv4-addr[source-ipv4-addr,interval,min-
rx,multiplier,dest-ipv4-prefix].
    - Example: --ipv4-static-route-bfd 10.25.25.100,1.1.1.0,123,456,3,10.10.10.1/32

- IPv6 Static Route BFD Configuration:
    - The --ipv6-static-route-bfd option is used to configure IPv6 static route BFD.
    - The format for SLXOS is: device-ip,dest-ipv6-addr,source-ipv6-addr[interval,min-
rx,multiplier].
    - Example: --ipv6-static-route-bfd 10.25.25.100,1::1,2::2,300,300,3
    - The format for ONEOS is: device-ip,dest-ipv6-addr[source-ipv6-addr,interval,min-
rx,multiplier,dest-ipv6-prefix].
    - Example: --ipv6-static-route-bfd 10.25.25.100,1::1,,300,300,3,f::f/128
```

Note

- BFD source is applicable only for SLX-OS.
- BFD destination prefix is applicable only for OS ONE.

VRF Static Route BFD REST Model

The REST model for VRF static route BFD has the following attributes:

- BFD Source IP: An optional attribute applicable only for SLX-OS and ignored forOS ONE.
- Destination Prefix: An optional attribute applicable only for OS ONE and ignored for SLX-OS.

Backward Compatibility

The following examples demonstrate backward compatibility for VRF create and update operations.

- VRF Create

```
$efa tenant vrf create --name redslx --tenant t1_slx
            --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.41
            --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.42
            --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.43
Vrf created successfully.
```

To view the running-config of all current XCO configurations for core services, run the **`show-running-config`** command.

```
SLX-69# show running-config vrf redslx
vrf redslx
 rd 172.31.254.1:1
```

```
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
  ip route 10.10.10.1/32 10.20.30.41
  ip route 10.10.10.1/32 10.20.30.42
  ip route 10.10.10.1/32 10.20.30.43
```

- VRF Update

```
$efa tenant vrf update --name redslx --tenant t1_slx
            --operation static-route-bfd-add
            --ipv4-static-route-bfd 10.64.196.69,10.20.30.41,1.1.1.0,123,456,3
            --ipv4-static-route-bfd 10.64.196.69,10.20.30.42,2.2.2.0,123,456,3
Vrf updated successfully.
```

To view the VRF configuration details, run the **`efa tenant vrf show`** command.

```
efa tenant vrf show --name redslx --tenant t1_slx --detail
================================================================================
Name                                  : redslx
Tenant                                : t1_slx
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              : false
Enable Default Information Originate    : false
Route Target                  :
Static Route        : Switch-IP->{Network,Nexthop-IP[Route-Distance,Route-Metric]}, ...
                : 10.64.196.69->{10.10.10.1/32,10.20.30.41[ , ]}
                                              {10.10.10.1/32,10.20.30.42[ , ]}
                                              {10.10.10.1/32,10.20.30.43[ , ]}
Static Route BFD    : Switch-IP->{DestIP,SourceIP[Interval,Min-
Rx,Multiplier,DestPrefix]}, ...
                [Flag: * - Derived Static Route BFD]
                : 10.64.196.69->{10.20.30.41,1.1.1.0[123,456,3,]}
                                              {10.20.30.42,2.2.2.0[123,456,3,]}
Network Route Address            :
Static Network            :
Aggregate Address             :
VRF Type                    : private
State                     : vrf-device-created
Dev State                 : provisioned
App State                 : cfg-in-sync

================================================================================
```

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

```
SLX-69# show running-config vrf redslx
vrf redslx
 rd 172.31.254.1:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
```

```
route-target export 101:101 evpn
route-target import 101:101 evpn
ip route static bfd 10.20.30.41 1.1.1.0 interval 123 min-rx 456 multiplier 3
ip route static bfd 10.20.30.42 2.2.2.0 interval 123 min-rx 456 multiplier 3
ip route 10.10.10.1/32 10.20.30.41
ip route 10.10.10.1/32 10.20.30.42
ip route 10.10.10.1/32 10.20.30.43
```

VRF Static Route BFD Provisioning for SLX-OS and OS ONE

Use this topic to learn about provisioning BFD for VRF static route on SLX-OS and OS ONE. The uniform provisioning model allows for consistent configuration of BFD for VRF static route across SLX-OS and OS ONE devices.

Configure BFD for VRF Static Route (Ignoring Source BFD for OS ONE) Using Current Approach

You can configure a static route and static route BFD using the current approach.

About This Task

Follow this procedure to configure static route and static route BFD on SLX-OS and OS ONE using the current approach.

Procedure

1. For SLX-OS, configure a static route before configuring BFD for the static route.

```
efa tenant vrf create --name redslx --tenant t1_slx
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.41
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.42
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.43
# Vrf created successfully
```

To view the running-config of all current XCO configurations for core services, run the **show-running-config** command.

```
SLX-69# show running-config vrf redslx
vrf redslx
 rd 172.31.254.1:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
  ip route 10.10.10.1/32 10.20.30.41
  ip route 10.10.10.1/32 10.20.30.42
  ip route 10.10.10.1/32 10.20.30.43
 !
```

```
efa tenant vrf update --name redslx --tenant t1_slx
    --operation static-route-bfd-add
    --ipv4-static-route-bfd 10.64.196.69,10.20.30.41,1.1.1.0,123,456,3
    --ipv4-static-route-bfd 10.64.196.69,10.20.30.42,2.2.2.0,123,456,3
# Vrf updated successfully
```

VRF Show Output:

- VRF redslx created with tenant t1_slx.
- Static routes configured at switch IP 10.64.196.69 pointing to networks via next hops.

- Static Route BFD entries show association with destinations 10.20.30.41 and 10.20.30.42.
- VRF state: created, not yet provisioned, application state ready.

```
efa tenant vrf show --name redslx --tenant t1_slx --detail
================================================================================
Name                                        : redslx
Tenant                                     : t1_slx
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                   : false
Enable Default Information Originate         : false
Route Target                    :
Static Route         : Switch-IP->{Network,Nexthop-IP[Route-Distance,Route-Metric]}, ...
                 : 10.64.196.69->{10.10.10.1/32,10.20.30.41[ , ]}
                                      {10.10.10.1/32,10.20.30.42[ , ]}
                                      {10.10.10.1/32,10.20.30.43[ , ]}
Static Route BFD    : Switch-IP->{DestIP,SourceIP[Interval,Min-
Rx,Multiplier,DestPrefix]}, ...
                 [Flag: * - Derived Static Route BFD]
                 : 10.64.196.69->{10.20.30.41,1.1.1.0[123,456,3,]}
                                      {10.20.30.42,2.2.2.0[123,456,3,]}
Network Route Address           :
Static Network              :
Aggregate Address              :
VRF Type                      : private
State                       : vrf-device-created
Dev State                  : provisioned
App State                  : cfg-in-sync
================================================================================
```

To view the running-config of all current XCO configurations for core services, run the **show-running-config** command.

```
SLX-69# show running-config vrf redslx
vrf redslx
 rd 172.31.254.1:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
  ip route static bfd 10.20.30.41 1.1.1.0 interval 123 min-rx 456 multiplier 3
  ip route static bfd 10.20.30.42 2.2.2.0 interval 123 min-rx 456 multiplier 3
  ip route 10.10.10.1/32 10.20.30.41
  ip route 10.10.10.1/32 10.20.30.42
  ip route 10.10.10.1/32 10.20.30.43
 !
```

2. For SLX, configure a static route BFD before configuring the static route.

```
efa tenant vrf create --name redslx --tenant t1_slx
    --ipv4-static-route-bfd 10.64.196.69,10.20.30.41,1.1.1.0,123,456,3
```

```
    --ipv4-static-route-bfd 10.64.196.69,10.20.30.42,2.2.2.0,123,456,3
# Vrf created successfully

efa tenant vrf update --name redslx --tenant t1_slx
    --operation static-route-add
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.41
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.42
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.43
# Vrf updated successfully
```

The VRF show output is similar to the output given in step one. Static routes and static route BFD entries are configured, but in the reversed order of commands compared to the previous step (step 1).

3. For OS ONE, configure a static route before configuring BFD for the static route.

```
efa tenant vrf create --name red --tenant t4
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.41
    --ipv4-static-route-next-hop 10.64.188.68,10.10.20.1/32,10.20.30.41
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.42
    --ipv4-static-route-next-hop 10.64.188.68,10.10.20.1/32,10.20.30.42
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.43
# Vrf created successfully

efa tenant vrf update --name red --tenant t4
    --operation static-route-bfd-add
    --ipv4-static-route-bfd 10.64.188.68,10.20.30.41,1.1.1.0,123,456,3
    --ipv4-static-route-bfd 10.64.188.68,10.20.30.42,2.2.2.0,123,456,3
# Vrf updated successfully
```

VRF Show Output:

- VRF red under tenant t4 configured with multiple static routes.
- Static Route BFD entries appear after update.
- Note that OS ONE ignores the source BFD in the current format.
- SR BFD entries related to SR are pushed to devices during Endpoint Group (EPG) creation or update.

4. For OS ONE, configure a static route BFD before configuring the static route.

```
efa tenant vrf create --name red --tenant t4 \
    --ipv4-static-route-bfd 10.64.188.68,10.20.30.41,1.1.1.0,123,456,3 \
    --ipv4-static-route-bfd 10.64.188.68,10.20.30.42,2.2.2.0,123,456,3
# Vrf created successfully


efa tenant vrf update --name red --tenant t4 \
    --operation static-route-add \
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.41 \
    --ipv4-static-route-next-hop 10.64.188.68,10.10.20.1/32,10.20.30.41 \
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.42 \
    --ipv4-static-route-next-hop 10.64.188.68,10.10.20.1/32,10.20.30.42 \
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.43
# Vrf updated successfully
```

VRF Show Output:

- Similar to above, with static-route BFD configured first followed by static routes.
- VRF red with tenant t4 shows respective routes and BFD configurations.
- For OS ONE, the source BFD is ignored under the current provisioning model. Static Route BFD entries associated with Static Routes are pushed to the device during EPG create or update operations.

Configure BFD for VRF Static Route (Ignoring Destination IP Prefix on SLX-OS) Using Current SLX Format and Destination Prefix

You can configure a static route and static route BFD using the current SLX format and Destination prefix.

About This Task

When configuring SLX-OS, the destination IP prefix within the static route commands is ignored.

Follow this procedure to configure static route and static route BFD on SLX-OS and OS ONE.

Procedure

1. For SLX, configure a static route before configuring BFD for the static route.

```
efa tenant vrf create --name redslx --tenant t1_slx
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.41
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.42
    --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.43
Vrf created successfully.
```

   - Run the **`efa tenant vrf show --<vrf name> --<tenant name> --detail`** command to view detailed VRF information showing static routes configured.

```
efa tenant vrf show --name redslx --tenant t1_slx --detail
================================================================================
Name                                    : redslx
Tenant                                 : t1_slx
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                : false
Enable Default Information Originate        : false
Route Target                :
Static Route        : Switch-IP->{Network,Nexthop-IP[Route-Distance,Route-
Metric]}, ...
               : 10.64.196.69->{10.10.10.1/32,10.20.30.41[ , ]}
                                        {10.10.10.1/32,10.20.30.42[ , ]}
                                        {10.10.10.1/32,10.20.30.43[ , ]}
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
================================================================================

SLX-69# show running-config vrf redslx
vrf redslx
 rd 172.31.254.1:1
```

```
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
  ip route 10.10.10.1/32 10.20.30.41
  ip route 10.10.10.1/32 10.20.30.42
  ip route 10.10.10.1/32 10.20.30.43
```

- Update the VRF to add static route BFD configuration:

```
efa tenant vrf update --name redslx --tenant t1_slx
           --operation static-route-bfd-add
           --ipv4-static-route-bfd
10.64.196.69,10.20.30.41,1.1.1.0,123,456,3,10.10.10.1/32
           --ipv4-static-route-bfd
10.64.196.69,10.20.30.42,2.2.2.0,123,456,3,10.10.10.1/32
```

- Confirm with the detailed VRF view that shows both static routes and SR BFD entries.

2. For SLX, configure a static route BFD before configuring the static route.

```
efa tenant vrf create --name redslx --tenant t1_slx
          --ipv4-static-route-bfd
10.64.196.69,10.20.30.41,1.1.1.0,123,456,3,10.10.10.1/32
          --ipv4-static-route-bfd
10.64.196.69,10.20.30.42,2.2.2.0,123,456,3,10.10.10.1/32
Vrf created successfully.

SLX-69# show running-config vrf redslx
vrf redslx
 rd 172.31.254.1:1
 evpn irb ve 8192 cluster-gateway
 address-family ipv4 unicast
  route-target export 101:101 evpn
  route-target import 101:101 evpn
  ip route static bfd 10.20.30.41 1.1.1.0 interval 123 min-rx 456 multiplier 3
  ip route static bfd 10.20.30.42 2.2.2.0 interval 123 min-rx 456 multiplier 3
 !

efa tenant vrf update --name redslx --tenant t1_slx
           --operation static-route-add
           --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.41
           --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.42
           --ipv4-static-route-next-hop 10.64.196.69,10.10.10.1/32,10.20.30.43
```

The detailed VRF display confirms both static route and BFD configurations.

3. For OS ONE, configure a static route before configuring BFD for the static route.

```
efa tenant vrf create --name red --tenant t4
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.41
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.42
    --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.43
Vrf created successfully.
```

- View VRF detail; static routes are present; no BFD yet.
- Update the VRF to add SR BFD.

```
efa tenant vrf update --name red --tenant t4
     --operation static-route-bfd-add
     --ipv4-static-route-bfd
10.64.188.68,10.20.30.41,1.1.1.0,123,456,3,10.10.10.1/32
     --ipv4-static-route-bfd
10.64.188.68,10.20.30.42,2.2.2.0,123,456,3,10.10.10.1/32
Vrf updated successfully.
```

- Confirm detailed view shows static route BFD entries tied to the static routes.

4. For OS ONE, configure a static route BFD before configuring the static route.

```
efa tenant vrf create --name red --tenant t4
     --ipv4-static-route-bfd 10.64.188.68,10.20.30.41,1.1.1.0,123,456,3,10.10.10.1/32
     --ipv4-static-route-bfd 10.64.188.68,10.20.30.42,2.2.2.0,123,456,3,10.10.10.1/32
Vrf created successfully.
```

- Detail shows only BFD entries initially.
- Update VRF to add static routes afterwards.

```
efa tenant vrf update --name red --tenant t4
     --operation static-route-add
     --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.41
     --ipv4-static-route-next-hop 10.64.188.68,10.10.10.1/32,10.20.30.42
Vrf updated successfully.
```

- Detailed VRF view confirms both static routes and associated BFD configurations.
- Static route BFD entries linked to SR routes are pushed to devices during Endpoint Group (EPG) create or update operations.

Note
Automatically derived BFD destination prefixes (auto-derived dest-prefix BFDs) cannot be deleted using the static-route-bfd-delete operation.

BGP Peer Group

Use this topic to learn about BGP peer group configuration and supported parameters on OS ONE.

Supported Parameters for Extreme OS ONE

#	BGP Peer Group Parameters	Support for Extreme OS ONE Device
1	`--pg-name`	**YES**
2	`--pg-asn`	**YES**
3	`--pg-bfd-enable`	**YES**
4	`--pg-bfd`	**YES**
5	`--pg-next-hop-self`	**YES**
6	`--pg-update-source-ip`	**YES**
7	`--pg-md5-password-prompt-enable`	**YES**
8	`--pg-md5-password`	**YES**
9	`--pg-remove-private-as`	NO
10	`--pg-ipv6-uc-nbr-activate`	**YES**
11	`--pg-ipv4-uc-nbr-prefix-list`	NO
12	`--pg-ipv6-uc-nbr-prefix-list`	NO
13	`--pg-ipv4-uc-nbr-route-map`	NO
14	`--pg-ipv6-uc-nbr-route-map`	NO
15	`--pg-ipv4-uc-nbr-send-community`	NO

#	BGP Peer Group Parameters	Support for Extreme OS ONE Device
16	--pg-ipv6-uc-nbr-send-community	NO
17	--pg-ipv4-uc-nbr-add-path-capability	**YES**
18	--pg-ipv6-uc-nbr-add-path-capability	**YES**
19	--pg-ipv4-uc-nbr-add-path-advertise-all	NO
20	--pg-ipv6-uc-nbr-add-path-advertise-all	NO
21	--pg-ipv4-uc-nbr-add-path-advertise-best	NO
22	--pg-ipv6-uc-nbr-add-path-advertise-best	NO
23	--pg-ipv4-uc-nbr-add-path-advertise-group-best	NO
24	--pg-ipv6-uc-nbr-add-path-advertise-group-best	NO

BGP Peer Group CLI Model

Use the **efa tenant service bgp peer-group create** command to create a BGP peer group instance.

```
efa tenant service bgp peer-group create --tenant t2 --name SLX1 --pg-name
10.64.160.11:pg1 --pg-asn 10.64.160.11,pg1:651000
```

BGP Peer Group REST Model

A new attribute "vrf-name" is added to the BGP peer group model, applicable for OS ONE and ignored for SLX-OS.

Table 38: Default-VRF vs Tenant-VRF

	Default-VRF	Tenant-VRF
Scope	System-level, fabric-wide	User-defined, tenant-specific
Purpose	Core routing, underlay, gateway connectivity	Tenant isolation, private routing
Use Cases	InterAS, GW handoff, BGP peering	Tenant segmentation, app isolation
Key Point	Globally consistent across fabric	Local to tenant, not shared

Default-VRF handles global (fabric-wide) routing whereas Tenant-VRF handles tenant-specific isolation and routing.

Backward Compatibility

The following examples demonstrate backward compatibility for BGP peer group create and show operations:

- BGP Peer Group Create

```
efa tenant service bgp peer-group create --tenant t2 --name SLX1
                   --pg-name 10.64.160.11:pg1
                 --pg-asn 10.64.160.11,pg1:651000
                   --pg-bfd-enable 10.64.160.11,pg1:true --pg-bfd
10.64.160.11,pg1:100,100,3
                   --pg-next-hop-self 10.64.160.11,pg1:true --pg-update-source-ip
10.64.160.11,pg1:10.10.10.10
                   --pg-md5-password
```

```
10.64.160.11,pg1:'$9$MCgKGaNt6OASX68/7TC6Lw=='
                  --pg-ipv4-uc-nbr-add-path-capability 10.64.160.11,pg1:receive

BgpService created successfully.
```

--pg-asn is a mandatory parameter for SLX and OS ONE.

- BGP Peer Group Show

  Use the **efa tenant service bgp peer-group show** command to display the BGP peer group configuration details, including device IP, peer group VRF name, peer group, remote ASN, activate, next hop self, update source IP, BFD enabled, and BFD interval, Rx, and multiplier.

```
efa tenant service bgp peer-group show --name SLX1 --tenant t2
================================================================================================
================================================================================
Name        : SLX1
Tenant      : t2
State       : bgp-pg-created

+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
|  Device IP   | Peer Group VRF Name | Peer Group | Remote |        Activate        |
Next Hop |  Update       |   BFD   |           BFD            |  Dev State  |  App
State  |
|              |                     |            |  ASN   |  [AFI,SAFI->Activate]
|   Self   | Source IP     | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
| 10.64.160.11 |    default-vrf      |    pg1     | 651000 | ipv4, unicast -> true  |
true     | 10.10.101.10  |  true   |  100, 100,3              | provisioned | cfg-in-
sync |
|              |                     |            |        | ipv6, unicast -> false
|          |               |         |                          |
|             |
|              |                     |            |        |
|          |               |         |                          |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
BGP PeerGroup Details
================================================================================================
================================================================================
--- Time Elapsed: 235.9579ms ---
```

- BGP Peer Group Show Detail

  Use this to display the BGP peer group configuration information, including description, device IP, peer group, peer group VRF name, remote ASN, next hop self, update source IP, BFD enabled, BFD interval, Rx, and multiplier, MD5 password, and dev state.

```
efa tenant service bgp peer-group show --name SLX1 --tenant t2 --detail
================================================================================================
================================================================================
Name             : SLX1
Tenant           : t2
State            : bgp-pg-created
Description      :
```

```
Peer Group
----------
        Device IP            : 10.64.160.11
        Peer Group           : pg1
        Peer Group VRF Name  : default-vrf
        Remote ASN           : 651000
        Next Hop Self        : true
        Update Source IP     : 10.10.10.10
        BFD Enabled          : true
        BFD Interval         : 100
        BFD Rx               : 100
        BFD Multiplier       : 3
        MD5 Password         : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  : Receive (ipv4)
        Dev State            : provisioned
        App State            : cfg-in-sync

================================================================================
================================================================================

router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 651000
 neighbor pg1 update-source 10.10.10.10
 neighbor pg1 next-hop-self
 neighbor pg1 password $9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-rx 100 multiplier 3
 neighbor 10.20.20.3 remote-as 4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.57/32
  network 172.31.254.107/32
  neighbor pg1 additional-paths receive
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
 !
```

BGP Peer Group Provisioning for SLX-OS and OS ONE

Use this topic to learn about provisioning BGP peer group on SLX-OS and OS ONE. The uniform provisioning model allows for consistent configuration of BGP peer group across SLX-OS and OS ONE devices.

Configure BGP Peer Group Without VRF Using Existing Format

You can configure a BGP peer group without a VRF using the existing format.

About This Task

Follow this procedure to create a BGP peer group without a VRF using the existing format on SLX and OS ONE.

Note
The "default-vrf" is specified as the peer-group VRF.

Procedure

1. Create a BGP peer group on SLX.

```
efa tenant service bgp peer-group create --tenant t2 --name SLX1
         --pg-name 10.64.160.11:pg1 --pg-asn 10.64.160.11,pg1:651000 --pg-bfd-enable
10.64.160.11,pg1:true
         --pg-bfd 10.64.160.11,pg1:100,100,3  --pg-next-hop-self 10.64.160.11,pg1:true
         --pg-update-source-ip 10.64.160.11,pg1:10.10.10.10
         --pg-md5-password 10.64.160.11,pg1:'$9$MCgKGaNt6OASX68/7TC6Lw=='
         --pg-ipv4-uc-nbr-add-path-capability 10.64.160.11,pg1:receive
Bgp Service created successfully.
```

2. Show BGP Peer Group on SLX.

```
efa tenant service bgp peer-group show --name SLX1 --tenant t2
================================================================================================
================================================================================
Name        : SLX1
Tenant      : t2
State       : bgp-pg-created

+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+--------------+
|  Device IP   | Peer Group VRF Name | Peer Group | Remote |        Activate        |
Next Hop |  Update       |   BFD   |           BFD            |  Dev State  |  App
State  |
|              |                    |            |  ASN   |  [AFI,SAFI->Activate]
|   Self   | Source IP     | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+--------------+
| 10.64.160.11 |    default-vrf      |    pg1     | 651000 | ipv4, unicast -> true  |
true     | 10.10.101.10  |  true   |  100, 100,3              | provisioned | cfg-in-
sync |
|              |                    |            |        | ipv6, unicast -> false
|          |               |         |                          |
|             |
|              |                    |            |        |
|          |               |         |                          |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+--------------+
BGP PeerGroup Details
================================================================================================
================================================================================
--- Time Elapsed: 235.9579ms ---
```

3. Create a BGP peer group on OS ONE.

Note
Unlike SLX, OS ONE can have only one address-family configured at a time, so XCO supports single AFI per PG for OS ONE.

```
efa tenant service bgp peer-group create --tenant t2 --name TOS1
          --pg-name 10.64.184.86:pg1 --pg-asn 10.64.184.86,pg1:651000 --pg-bfd-enable
10.64.184.86,pg1:true
          --pg-bfd 10.64.184.86,pg1:100,100,3  --pg-next-hop-self
10.64.184.86,pg1:true
          --pg-update-source-ip 10.64.184.86,pg1:10.10.10.10 --pg-md5-password
10.64.184.86,pg1:'password'
          --pg-ipv4-uc-nbr-add-path-capability 10.64.184.86,pg1:receive

BgpService created successfully.
```

4. Show BGP Peer Group on OS ONE.

```
efa tenant service bgp peer-group show --name TOS1 --tenant t2
================================================================================================
================================================================================
Name        : TOS1
Tenant      : t2
State       : bgp-pg-created

+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
|  Device IP   | Peer Group VRF Name | Peer Group | Remote |        Activate        |
Next Hop |  Update       |   BFD   |           BFD            |  Dev State  |  App
State  |
|              |                     |            |  ASN   |  [AFI,SAFI->Activate]
|   Self   | Source IP     | Enabled | [Interval,Rx,Multiplier] |
|              |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
| 10.64.184.86 |    default-vrf      |    pg1     | 651000 | ipv4, unicast -> true  |
true     | 10.10.101.10  |  true   |  100, 100,3              | provisioned | cfg-in-
sync |
|              |                     |            |        | ipv6, unicast -> false
|          |               |         |                          |
|             |
|              |                     |            |        |
|          |               |         |                          |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
BGP PeerGroup Details
================================================================================================
================================================================================
--- Time Elapsed: 235.9579ms ---

$ efa tenant service bgp peer-group show --name TOS1 --tenant t2 --detail
================================================================================================
================================================================================
Name             : TOS1
Tenant           : t2
State            : bgp-pg-created
Description      :

Peer Group
```

```
----------
        Device IP            : 10.64.184.86
        Peer Group           : pg1
        Peer Group VRF Name  : default-vrf
        Remote ASN           : 651000
        Next Hop Self        : true
        Update Source IP     : 10.10.10.10
        BFD Enabled          : true
        BFD Interval         : 100
        BFD Rx               : 100
        BFD Multiplier       : 3
        MD5 Password         : password
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  : Receive (ipv4)
        Dev State            : provisioned
        App State            : cfg-in-sync

================================================================================
================================================================================

vrf default-vrf
 .....
  router bgp
    local-as 4200000000
    router-id 172.31.254.63
    route-distinguisher 172.31.254.63:1
    …..
     peer-group pg1
      remote-as 651000
      auth-password-hashed
0b14d501a594442a01c6859541bcb3e8164d183d32937b851835442f69d5c94e
      enable-bfd profile profile_100_100_3
      update-source 10.10.10.10
      address-family ipv4 unicast
        nexthop-self
        add-paths receive
        activate
      !
    !
  !
!
```

Configure BGP Peer Group With VRF Using New Format

You can configure a BGP peer group with a VRF using the new format.

About This Task

The new format allows specifying VRF input per peer group. This input is ignored for SLX-OS and honored for OS ONE.

Follow this procedure to create a BGP peer group on SLX and OS ONE with the new format.

Procedure

1. Create a BGP peer group on SLX.

```
efa tenant service bgp peer-group create --tenant t2 --name SLX1 --pg-name
10.64.160.11:pg1;default-vrf --pg-asn 10.64.160.11,pg1,default-vrf:651000 --pg-bfd-
enable 10.64.160.11,pg1,default-vrf:true --pg-bfd 10.64.160.11,pg1,default-
vrf:100,100,3  --pg-next-hop-self 10.64.160.11,pg1,default-vrf:true --pg-update-source-
ip 10.64.160.11,pg1,default-vrf:10.10.10.10 --pg-md5-password 10.64.160.11,pg1,default-
vrf:'$9$MCgKGaNt6OASX68/7TC6Lw=='     --pg-ipv4-uc-nbr-add-path-capability
10.64.160.11,pg1,default-vrf:receive

BgpService created successfully.

efa tenant service bgp peer-group show --name SLX1 --tenant t2
================================================================================================
=========================================================================
Name        : SLX1
Tenant      : t2
State       : bgp-pg-created

+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
|  Device IP   | Peer Group VRF Name | Peer Group | Remote |        Activate        |
Next Hop |  Update       |   BFD   |           BFD            |  Dev State  |  App
State  |
|              |                    |            |  ASN   |  [AFI,SAFI->Activate]
|   Self   | Source IP     | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
| 10.64.160.11 |    default-vrf      |    pg1     | 651000 | ipv4, unicast -> true  |
true     | 10.10.101.10  |  true   |  100, 100,3              | provisioned | cfg-in-
sync |
|              |                    |            |        | ipv6, unicast -> false
|          |               |         |                          |
|             |
|              |                    |            |        |
|          |               |         |                          |
|             |
+--------------+--------------------+------------+--------+------------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
BGP PeerGroup Details
================================================================================================
=========================================================================
--- Time Elapsed: 235.9579ms ---

$ efa tenant service bgp peer-group show --name SLX1 --tenant t2 --detail
================================================================================================
=========================================================================
Name             : SLX1
Tenant           : t2
State            : bgp-pg-created
Description      :

Peer Group
----------
        Device IP            : 10.64.160.11
        Peer Group           : pg1
        Peer Group VRF Name  : default-vrf
        Remote ASN           : 651000
        Next Hop Self        : true
```

```
        Update Source IP     : 10.10.10.10
        BFD Enabled          : true
        BFD Interval         : 100
        BFD Rx               : 100
        BFD Multiplier       : 3
        MD5 Password         : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  : Receive (ipv4)
        Dev State            : provisioned
        App State            : cfg-in-sync

================================================================================
================================================================================

router bgp
 local-as 4200000000
 capability as4-enable
 fast-external-fallover
 neighbor pg1 peer-group
 neighbor pg1 remote-as 651000
 neighbor pg1 update-source 10.10.10.10
 neighbor pg1 next-hop-self
 neighbor pg1 password $9$MCgKGaNt6OASX68/7TC6Lw==
 neighbor pg1 bfd
 neighbor pg1 bfd interval 100 min-rx 100 multiplier 3
 neighbor 10.20.20.3 remote-as 4200000000
 neighbor 10.20.20.3 next-hop-self
 address-family ipv4 unicast
  network 172.31.254.57/32
  network 172.31.254.107/32
  neighbor pg1 additional-paths receive
  maximum-paths 8
  graceful-restart
 !
 address-family ipv6 unicast
 !
 address-family l2vpn evpn
  graceful-restart
```

2. Create a BGP peer group on OS ONE.

```
efa tenant service bgp peer-group create --tenant t2 --name TOS1 --pg-name
10.64.184.86:pg1;vrfRed --pg-asn 10.64.184.86,pg1:651000 --pg-bfd-enable
10.64.184.86,pg1:true --pg-bfd 10.64.184.86,pg1:100,100,3  --pg-next-hop-self
10.64.184.86,pg1:true --pg-update-source-ip 10.64.184.86,pg1:10.10.10.10 --pg-md5-
password 10.64.184.86,pg1:'password'     --pg-ipv4-uc-nbr-add-path-capability
10.64.184.86,pg1:receive

BgpService created successfully.

efa tenant service bgp peer-group show --name TOS1 --tenant t2
================================================================================
================================================================================
Name        : TOS1
Tenant      : t2
State       : bgp-pg-created
```

```
+--------------+--------------------+------------+--------+-----------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
|  Device IP   | Peer Group VRF Name | Peer Group | Remote |        Activate        |
Next Hop |  Update       |   BFD   |           BFD            |  Dev State  |  App
State  |
|              |                    |            |  ASN   |  [AFI,SAFI->Activate]
|   Self   | Source IP     | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+--------------------+------------+--------+-----------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
| 10.64.184.86 |    vrfRed           |    pg1     | 651000 | ipv4, unicast -> true  |
true     | 10.10.101.10  |  true   |  100, 100,3              | provisioned | cfg-in-
sync |
|              |                    |            |        | ipv6, unicast -> false
|          |               |         |                          |
|             |
|              |                    |            |        |
|          |               |         |                          |
|             |
+--------------+--------------------+------------+--------+-----------------------
+----------+---------------+---------+--------------------------+-------------
+-------------+
BGP PeerGroup Details
================================================================================================
=========================================================================
--- Time Elapsed: 235.9579ms ---

$ efa tenant service bgp peer-group show --name TOS1 --tenant t2 --detail
================================================================================================
=========================================================================
Name             : TOS1
Tenant           : t2
State            : bgp-pg-created
Description      :

Peer Group
----------
        Device IP            : 10.64.184.86
        Peer Group           : pg1
        Peer Group VRF Name  : vrfRed
        Remote ASN           : 651000
        Next Hop Self        : true
        Update Source IP     : 10.10.10.10
        BFD Enabled          : true
        BFD Interval         : 100
        BFD Rx               : 100
        BFD Multiplier       : 3
        MD5 Password         : password
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  : Receive (ipv4)
        Dev State            : provisioned
        App State            : cfg-in-sync

================================================================================================
```

```
=============================================================================


vrf vrfRed
 .....
  router bgp
    local-as 4200000000
    router-id 172.31.254.63
    route-distinguisher 172.31.254.63:1
    as-notation as-plain
    address-family ipv4 unicast
      use-multiple-paths ibgp maximum-paths 8
      use-multiple-paths ebgp maximum-paths 8
      activate
    !
    address-family ipv6 unicast
      use-multiple-paths ibgp maximum-paths 8
      use-multiple-paths ebgp maximum-paths 8
      activate
    !
    instance-type evpn
      nve fabs
        l3vni 8173
        address-family ipv4 unicast
          route-target export 101:101
          route-target import 101:101
        !
        address-family ipv6 unicast
          route-target export 101:101
          route-target import 101:101
        !
      !
    !
     peer-group pg1
      remote-as 651000
      auth-password-hashed
0b14d501a594442a01c6859541bcb3e8164d183d32937b851835442f69d5c94e
      enable-bfd profile profile_100_100_3
      update-source 10.10.10.10
      address-family ipv4 unicast
        nexthop-self
        add-paths receive
        activate
```

Upgrade Handling

During the upgrade, existing peer groups are automatically associated with the default VRF, and the server app state gets updated for the default VRF.

Before Upgrade

The BGP peer group configuration is updated and displayed using the **`efa tenant service bgp peer-group show`** command.

- Show BGP Peer Group

```
efa tenant service bgp peer-group show --tenant p2
================================================================================================
=============================================================================
Name        : p1
Tenant      : p2
State       : bgp-pg-created
```

```
+--------------+--------------------+------------+--------+-------------------------
+----------+-------------+---------+--------------------------+-------------
+-------------+
|  Device IP   | Peer Group VRF Name | Peer Group | Remote |        Activate        |
Next Hop |   Update    |   BFD   |           BFD            |  Dev State  |  App
State  |
|              |                    |            |  ASN   |  [AFI,SAFI->Activate]
|   Self   |  Source IP  | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+--------------------+------------+--------+-------------------------
+----------+-------------+---------+--------------------------+-------------
+-------------+
| 10.64.196.61 |                    |    pg1     | 651000 | ipv4, unicast -> true
|   true   | 10.10.10.10 |  true   |       100, 100, 3        | provisioned | cfg-in-
sync |
|              |                    |            |        | ipv6, unicast -> false
|          |             |         |                          |
|             |
+--------------+--------------------+------------+--------+-------------------------
+----------+-------------+---------+--------------------------+-------------
+-------------+
| 10.64.196.61 |                    |    pg2     | 651000 | ipv4, unicast -> true
|   true   | 10.10.10.10 |  true   |       100, 100, 3        | provisioned | cfg-in-
sync |
|              |                    |            |        | ipv6, unicast -> false
|          |             |         |                          |
|             |
+--------------+--------------------+------------+--------+-------------------------
+----------+-------------+---------+--------------------------+-------------
+-------------+
BGP PeerGroup Details

================================================================================================
=========================================================================

--- Time Elapsed: 93.914107ms ---
```

After Upgrade

The BGP peer group configuration is displayed after the upgrade, showing the association with the default VRF.

```
efa tenant service bgp peer-group show --tenant p2
================================================================================================
=========================================================================
Name        : p1
Tenant      : p2
State       : bgp-pg-created

+--------------+-------------+------------+--------+-------------------------+----------
+-------------+---------+--------------------------+-------------+---------------+
|  Device IP   |  VRF Name   | Peer Group | Remote |        Activate         | Next Hop
|   Update    |   BFD   |           BFD            |  Dev State  |   App State   |
|              |             |            |  ASN   |  [AFI,SAFI->Activate]  |   Self   |
Source IP  | Enabled | [Interval,Rx,Multiplier] |             |               |
+--------------+-------------+------------+--------+-------------------------+----------
+-------------+---------+--------------------------+-------------+---------------+
| 10.64.196.61 | default-vrf |    pg1     | 651000 | ipv4, unicast -> true  |   true   |
10.10.10.10 |  true   |       100, 100, 3        | provisioned |  cfg-in-sync  |
|              |             |            |        | ipv6, unicast -> false |
|             |         |                          |             |               |
+--------------+-------------+------------+--------+-------------------------+----------
+-------------+---------+--------------------------+-------------+---------------+
| 10.64.196.61 | default-vrf |    pg2     | 651000 | ipv4, unicast -> true  |   true   |
```

```
10.10.10.10 |  true   |       100, 100, 3        | provisioned | cfg-in-sync |
|             |         |           |              | ipv6, unicast -> false |
|             |         |                        |             |              |
+-------------+---------+-----------+--------+-----------------------+----------
+-------------+---------+-------------------------+-------------+---------------+
BGP PeerGroup Details

================================================================================================
=====================================================================

--- Time Elapsed: 73.883898ms ---
(efa:user)root@dev-vm-11-14:/home/user/tnair/efa# efa tenant service bgp peer-group show
--tenant p2 --detail
================================================================================================
=====================================================================
Name             : p1
Tenant           : p2
State            : bgp-pg-created
Description      :

Peer Group
----------
        Device IP            : 10.64.196.61
        VRF Name             : default-vrf
        Peer Group           : pg1
        Remote ASN           : 651000
        Next Hop Self        : true
        Update Source IP     : 10.10.10.10
        BFD Enabled          : true
        BFD Interval         : 100
        BFD Rx               : 100
        BFD Multiplier       : 3
        MD5 Password         : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
        Route Map Out        :
        Send Community       :
        Add Path Capability  : Receive (ipv4)
        Dev State            : provisioned
        App State            : cfg-in-sync

        Device IP            : 10.64.196.61
        VRF Name             : default-vrf
        Peer Group           : pg2
        Remote ASN           : 651000
        Next Hop Self        : true
        Update Source IP     : 10.10.10.10
        BFD Enabled          : true
        BFD Interval         : 100
        BFD Rx               : 100
        BFD Multiplier       : 3
        MD5 Password         : $9$MCgKGaNt6OASX68/7TC6Lw==
        Remove Private AS    : false
        Activate             : Activate (afi)
                               true (ipv4)
                               false (ipv6)
        Prefix List In       :
        Prefix List Out      :
        Route Map In         :
```

```
        Route Map Out         :
        Send Community        :
        Add Path Capability   : Receive (ipv4)
        Dev State             : provisioned
        App State             : cfg-refreshed

================================================================================================
================================================================
```

BGP Static Peer

Use this topic to learn about BGP static peer configuration and supported parameters on OS ONE.

Supported Parameters for Extreme OS ONE

#	BGP Peer Parameters	Support for Extreme OS ONE Device
1	`--ipv4-uc-nbr`	**YES**
2	`--ipv6-uc-nbr`	**YES**
3	`--md5-password-prompt-enable`	NO
4	`--ipv4-uc-nbr-md5-password`	NO
5	`--ipv6-uc-nbr-md5-password`	NO
6	`--ipv4-uc-nbr-next-hop-self`	NO
7	`--ipv6-uc-nbr-next-hop-self`	NO
8	`--ipv4-uc-nbr-bfd`	**YES**
9	`--ipv6-uc-nbr-bfd`	**YES**
10	`--ipv4-uc-nbr-update-source-ip`	**YES**
11	`--ipv6-uc-nbr-update-source-ip`	**YES**
12	`--ipv4-uc-nbr-remove-private-as`	N
13	`--ipv6-uc-nbr-remove-private-as`	NO
14	`--ipv4-uc-nbr-default-originate`	NO
15	`--ipv6-uc-nbr-default-originate`	NO
16	`--ipv4-uc-nbr-default-originate-route-map`	NO
17	`--ipv6-uc-nbr-default-originate-route-map`	NO
18	`--ipv4-uc-nbr-prefix-list`	NO
19	`--ipv6-uc-nbr-prefix-list`	NO
20	`--ipv4-uc-nbr-route-map`	NO
21	`--ipv6-uc-nbr-route-map`	NO
22	`--ipv4-uc-nbr-send-community`	NO
23	`--ipv6-uc-nbr-send-community`	NO
24	`--ipv4-uc-nbr-add-path-capability`	NO

#	BGP Peer Parameters	Support for Extreme OS ONE Device
25	--ipv6-uc-nbr-add-path-capability	NO
26	--ipv4-uc-nbr-add-path-advertise-all	NO
27	--ipv6-uc-nbr-add-path-advertise-all	NO
28	--ipv4-uc-nbr-add-path-advertise-group-best	NO
29	--ipv6-uc-nbr-add-path-advertise-group-best	NO
30	--ipv4-uc-nbr-activate-in-ipv6-uc-af	NO
31	--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af	NO
32	--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af	NO
33	--ipv4-uc-nbr-route-map-in-ipv6-uc-af	NO
34	--ipv4-uc-nbr-send-community-in-ipv6-uc-af	NO
35	--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af	NO
36	--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af	NO
37	--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af	NO
38	--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af	NO
39	--ipv4-uc-dyn-nbr	**YES**
40	--ipv6-uc-dyn-nbr	**YES**

BGP Static Peer REST Model

A new optional attribute "peer-group-name" is introduced in the BGP static peer REST model, applicable for both SLX-OS and OS ONE.

Backward Compatibility

The following example demonstrates backward compatibility for BGP static peer creation.

- Create BGP Static Peer

```
efa tenant service bgp peer create --name
B10 --tenant t1 --ipv4-uc-nbr 10.64.196.53,vrf11:1.1.1.11,65001
--ipv4-uc-nbr-bfd 10.64.196.53,vrf11:1.1.1.11,true,50,5000,50 --ipv4-
uc-nbr-update-source-ip 10.64.196.53,vrf11:1.1.1.11,10.11.12.13
--ipv4-uc-dyn-nbr 10.64.196.53,vrf11:10.0.0.0/29,pg1,8 --ipv6-uc-
nbr 10.64.196.53,vrf11:2000:db8:1::,65001 --ipv6-uc-nbr-bfd
10.64.196.53,vrf11:2000:db8:1::,true,50,5000,50 --ipv6-uc-nbr-update-source-
ip 10.64.196.53,vrf11:2000:db8:1::,ff::ff --ipv6-uc-dyn-nbr
10.64.196.53,vrf11:8001:db8:1::/68,pg2

BGP Peer created successfully.
--- Time Elapsed: 5.82378375s ---
```

- Show BGP Peer

```
efa tenant service bgp peer show --tenant t1
================================================================================
================================================================================
==============
Name        : B10
Tenant      : t1
State       : bgp-peer-created
```

```
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
|  Device IP   |  VRF  | PeerGroup | AFI  |  SAFI   |  Remote IP   | Remote ASN |
Activate | Next Hop |   Update    |   BFD   |           BFD            |  Dev State
|  App State  |
|              |       |           |      |         |              |
|          |   Self   |  Source IP  | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
| 10.64.196.53 | vrf11 |           | ipv6 | unicast | 2000:db8:1:: |   65001    |
true   |  false   |   ff::ff    |  true   |       50, 5000, 50       | provisioned |
cfg-in-sync |
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
| 10.64.196.53 | vrf11 |           | ipv4 | unicast |   1.1.1.11   |   65001    |
true   |  false   | 10.11.12.13 |  true   |       50, 5000, 50       | provisioned |
cfg-in-sync |
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
Static Peer Details

[Flag '*' indicates Multi protocol capability]


+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
|  Device-IP   |  VRF  | AFI  |  SAFI   |  Listen Range   | Listen | Peer Group |  Dev
State  |  App State  |
|              |       |      |         |                 | Limit  |
|             |             |
+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
| 10.64.196.53 | vrf11 | ipv6 | unicast | 8001:db8:1::/68 |        |    pg2     |
provisioned | cfg-in-sync |
+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
| 10.64.196.53 | vrf11 | ipv4 | unicast |   10.0.0.0/29   |   8    |    pg1     |
provisioned | cfg-in-sync |
+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
Dynamic Peer Details

================================================================================
================================================================================
==============

--- Time Elapsed: 88.680441ms ---
efa tenant service bgp peer show --tenant t1 --detail
================================================================================
================================================================================
==============
Name             : B10
Tenant           : t1
State            : bgp-peer-created
Description      :

Static Peer
-----------
```

```
        Device IP                   : 10.64.196.53
        VRF                         : vrf11
        Peer Group                  :
        AFI                         : ipv6
        SAFI                        : unicast
        Remote IP                   : 2000:db8:1::
        Remote ASN                  : 65001
        Activate                    : true
        Next Hop Self               : false
        Update Source IP            : ff::ff
        BFD Enabled                 : true
        BFD Interval                : 50
        BFD Rx                      : 5000
        BFD Multiplier              : 50
        MD5 Password                :
        Remove Private AS           : false
        Default Originate           : false
        Default Originate Route Map :
        Send Community              :
        Prefix List In              :
        Prefix List Out             :
        Route Map In                :
        Route Map Out               :
        Add Path Capability         :
        Add Path Advertise          :
        Dev State                   : provisioned
        App State                   : cfg-in-sync

        Device IP                   : 10.64.196.53
        VRF                         : vrf11
        Peer Group                  :
        AFI                         : ipv4
        SAFI                        : unicast
        Remote IP                   : 1.1.1.11
        Remote ASN                  : 65001
        Activate                    : true
        Next Hop Self               : false
        Update Source IP            : 10.11.12.13
        BFD Enabled                 : true
        BFD Interval                : 50
        BFD Rx                      : 5000
        BFD Multiplier              : 50
        MD5 Password                :
        Remove Private AS           : false
        Default Originate           : false
        Default Originate Route Map :
        Send Community              :
        Prefix List In              :
        Prefix List Out             :
        Route Map In                :
        Route Map Out               :
        Add Path Capability         :
        Add Path Advertise          :
        Dev State                   : provisioned
        App State                   : cfg-in-sync


Dynamic Peer
------------
        Device IP                   : 10.64.196.53
        VRF                         : vrf11
        AFI                         : ipv6
        SAFI                        : unicast
        Listen Range                : 8001:db8:1::/68
```

```
        Listen Limit                :
        Peer Group                  : pg2
        Dev State                   : provisioned
        App State                   : cfg-in-sync

        Device IP                   : 10.64.196.53
        VRF                         : vrf11
        AFI                         : ipv4
        SAFI                        : unicast
        Listen Range                : 10.0.0.0/29
        Listen Limit                : 8
        Peer Group                  : pg1
        Dev State                   : provisioned
        App State                   : cfg-in-sync

================================================================================
================================================================================
==============

--- Time Elapsed: 83.766866ms ---
vslx53# show running-config router bgp
router bgp
local-as 4200065535
capability as4-enable
fast-external-fallover
neighbor pg1 peer-group
neighbor pg1 remote-as 651000
neighbor pg1 update-source 10.10.10.10
neighbor pg1 next-hop-self
neighbor pg1 password $9$MCgKGaNt6OASX68/7TC6Lw==
neighbor pg1 bfd
neighbor pg1 bfd interval 100 min-rx 100 multiplier 3
neighbor pg2 peer-group
neighbor pg2 remote-as 651000
neighbor pg2 update-source 10.10.10.10
neighbor pg2 next-hop-self
neighbor pg2 password $9$MCgKGaNt6OASX68/7TC6Lw==
neighbor pg2 bfd
neighbor pg2 bfd interval 100 min-rx 100 multiplier 3
neighbor 10.20.20.3 remote-as 4200065535
neighbor 10.20.20.3 next-hop-self
address-family ipv4 unicast
  network 172.31.254.95/32
  network 172.31.254.233/32
  neighbor pg1 additional-paths receive
  maximum-paths 8
  graceful-restart
!
address-family ipv4 unicast vrf vrf11
  redistribute connected
  listen-range 10.0.0.0/29 peer-group pg1 limit 8
  neighbor 1.1.1.11 remote-as 65001
  neighbor 1.1.1.11 bfd
  neighbor 1.1.1.11 bfd interval 50 min-rx 5000 multiplier 50
  neighbor 1.1.1.11 update-source 10.11.12.13
  maximum-paths 8
!
address-family ipv6 unicast
  neighbor pg2 activate
  neighbor pg2 additional-paths receive
!
address-family ipv6 unicast vrf vrf11
  redistribute connected
  listen-range 8001:db8:1::/68 peer-group pg2
```

```
  neighbor 2000:db8:1:: remote-as 65001
  neighbor 2000:db8:1:: bfd
  neighbor 2000:db8:1:: bfd interval 50 min-rx 5000 multiplier 50
  neighbor 2000:db8:1:: update-source ff::ff
  neighbor 2000:db8:1:: activate
  maximum-paths 8
!
address-family l2vpn evpn
  graceful-restart

address-family ipv4 unicast vrf vrf11
  redistribute connected
  listen-range 10.0.0.0/29 peer-group pg1 limit 8
  neighbor 1.1.1.11 remote-as 65001
  neighbor 1.1.1.11 bfd
  neighbor 1.1.1.11 bfd interval 50 min-rx 5000 multiplier 50
  neighbor 1.1.1.11 update-source 10.11.12.13
  maximum-paths 8
!
address-family ipv6 unicast
  neighbor pg2 activate
  neighbor pg2 additional-paths receive
!
address-family ipv6 unicast vrf vrf11
  redistribute connected
  listen-range 8001:db8:1::/68 peer-group pg2
  neighbor 2000:db8:1:: remote-as 65001
  neighbor 2000:db8:1:: bfd
  neighbor 2000:db8:1:: bfd interval 50 min-rx 5000 multiplier 50
  neighbor 2000:db8:1:: update-source ff::ff
  neighbor 2000:db8:1:: activate
  maximum-paths 8
!
address-family l2vpn evpn
  graceful-restart
!
!
```

BGP Static Peer Provisioning for SLX-OS and OS ONE

Use this topic to learn about provisioning BGP static peer on SLX-OS and OS ONE. The uniform provisioning model allows for consistent configuration of BGP static peers across SLX-OS and OS ONE devices.

BGP Static Peer Configuration Without Peer Group Using Existing Format

The existing format does not support peer group association.

- **SLX**: Backward compatibility is maintained, as described in previous topics.
- **OS ONE**: This format is not supported.

QoS Limitations in Extreme OS ONE

- Only QoS map and trust value configurations are supported
- No support for drop precedence, traffic shaping, or bandwidth allocation

Configure BGP Static Peer With Peer Group Using New Format

You can configure a BGP static peer with peer group input using the new format.

About This Task

The new format allows specifying a peer group for association.

Procedure

1. Create a BGP peer on SLX-OS.

```
efa tenant service bgp peer create --name
B12 --tenant t1 --ipv4-uc-nbr 10.64.196.53,vrf11:1.1.1.11,pg1
--ipv4-uc-nbr-bfd 10.64.196.53,vrf11:1.1.1.11,true,50,5000,50 --ipv4-uc-nbr-
update-source-ip 10.64.196.53,vrf11:1.1.1.11,10.11.12.13  --ipv4-uc-dyn-nbr
10.64.196.53,vrf11:10.0.0.0/29,pg1,8 --ipv6-uc-nbr 10.64.196.53,vrf11:2000:db8:1::,pg2
--ipv6-uc-nbr-bfd 10.64.196.53,vrf11:2000:db8:1::,true,50,5000,50 --ipv6-uc-nbr-
update-source-ip 10.64.196.53,vrf11:2000:db8:1::,ff::ff --ipv6-uc-dyn-nbr
10.64.196.53,vrf11:8001:db8:1::/68,pg2

BGP Peer created successfully.

--- Time Elapsed: 6.92792481s --
```

2. Show BGP peer on SLX-OS.

```
efa tenant service bgp peer show --tenant t1
================================================================================================
================================================================================================
==============
Name        : B12
Tenant      : t1
State       : bgp-peer-created

+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
|  Device IP   |  VRF  | PeerGroup | AFI  |  SAFI   |  Remote IP   | Remote ASN |
Activate | Next Hop |   Update    |   BFD   |           BFD            |  Dev State
|  App State  |
|              |       |           |      |         |              |
|           |   Self   |  Source IP  | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
| 10.64.196.53 | vrf11 |    pg2    | ipv6 | unicast | 2000:db8:1:: |            |
true   |  false   |   ff::ff    |  true   |       50, 5000, 50       | provisioned |
cfg-in-sync |
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
| 10.64.196.53 | vrf11 |    pg1    | ipv4 | unicast |   1.1.1.11   |            |
true   |  false   | 10.11.12.13 |  true   |       50, 5000, 50       | provisioned |
cfg-in-sync |
+--------------+-------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+--------------------------+-------------
+-------------+
Static Peer Details

[Flag '*' indicates Multi protocol capability]


+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
```

```
|  Device-IP   |  VRF  | AFI  |  SAFI   |  Listen Range   | Listen | Peer Group |  Dev
State  |  App State  |
|              |       |      |         |                 | Limit  |
|             |             |
+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
| 10.64.196.53 | vrf11 | ipv6 | unicast | 8001:db8:1::/68 |        |    pg2     |
provisioned | cfg-in-sync |
+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
| 10.64.196.53 | vrf11 | ipv4 | unicast |   10.0.0.0/29   |   8    |    pg1     |
provisioned | cfg-in-sync |
+--------------+-------+------+---------+-----------------+--------+------------
+-------------+-------------+
Dynamic Peer Details

================================================================================================
================================================================================================
==============

--- Time Elapsed: 67.287818ms ---
mkulkarni@EXT-6XTL334:~/GoDCApp/XCO/GoCommon/bin$ ./efa tenant service bgp peer show --
tenant t1 --detail
================================================================================================
================================================================================================
==============
Name             : B12
Tenant           : t1
State            : bgp-peer-created
Description      :

Static Peer
-----------
        Device IP                   : 10.64.196.53
        VRF                         : vrf11
        Peer Group                  : pg1
        AFI                         : ipv4
        SAFI                        : unicast
        Remote IP                   : 1.1.1.11
        Remote ASN                  :
        Activate                    : true
        Next Hop Self               : false
        Update Source IP            : 10.11.12.13
        BFD Enabled                 : true
        BFD Interval                : 50
        BFD Rx                      : 5000
        BFD Multiplier              : 50
        MD5 Password                :
        Remove Private AS           : false
        Default Originate           : false
        Default Originate Route Map :
        Send Community              :
        Prefix List In              :
        Prefix List Out             :
        Route Map In                :
        Route Map Out               :
        Add Path Capability         :
        Add Path Advertise          :
        Dev State                   : provisioned
        App State                   : cfg-in-sync

        Device IP                   : 10.64.196.53
        VRF                         : vrf11
        Peer Group                  : pg2
```

```
        AFI                          : ipv6
        SAFI                         : unicast
        Remote IP                    : 2000:db8:1::
        Remote ASN                   :
        Activate                     : true
        Next Hop Self                : false
        Update Source IP             : ff::ff
        BFD Enabled                  : true
        BFD Interval                 : 50
        BFD Rx                       : 5000
        BFD Multiplier               : 50
        MD5 Password                 :
        Remove Private AS            : false
        Default Originate            : false
        Default Originate Route Map :
        Send Community               :
        Prefix List In               :
        Prefix List Out              :
        Route Map In                 :
        Route Map Out                :
        Add Path Capability          :
        Add Path Advertise           :
        Dev State                    : provisioned
        App State                    : cfg-in-sync


Dynamic Peer
------------
        Device IP                    : 10.64.196.53
        VRF                          : vrf11
        AFI                          : ipv4
        SAFI                         : unicast
        Listen Range                 : 10.0.0.0/29
        Listen Limit                 : 8
        Peer Group                   : pg1
        Dev State                    : provisioned
        App State                    : cfg-in-sync

        Device IP                    : 10.64.196.53
        VRF                          : vrf11
        AFI                          : ipv6
        SAFI                         : unicast
        Listen Range                 : 8001:db8:1::/68
        Listen Limit                 :
        Peer Group                   : pg2
        Dev State                    : provisioned
        App State                    : cfg-in-sync

================================================================================
================================================================================
==============

--- Time Elapsed: 79.992195ms ---
vslx53# show running-config router bgp
router bgp
local-as 4200065535
capability as4-enable
fast-external-fallover
neighbor pg1 peer-group
neighbor pg1 remote-as 651000
neighbor pg1 update-source 10.10.10.10
neighbor pg1 next-hop-self
neighbor pg1 password $9$MCgKGaNt6OASX68/7TC6Lw==
neighbor pg1 bfd
```

```
neighbor pg1 bfd interval 100 min-rx 100 multiplier 3
neighbor pg2 peer-group
neighbor pg2 remote-as 651000
neighbor pg2 update-source 10.10.10.10
neighbor pg2 next-hop-self
neighbor pg2 password $9$MCgKGaNt6OASX68/7TC6Lw==
neighbor pg2 bfd
neighbor pg2 bfd interval 100 min-rx 100 multiplier 3
neighbor 10.20.20.3 remote-as 4200065535
neighbor 10.20.20.3 next-hop-self
address-family ipv4 unicast
  network 172.31.254.95/32
  network 172.31.254.233/32
  neighbor pg1 additional-paths receive
  maximum-paths 8
  graceful-restart
!
address-family ipv4 unicast vrf vrf11
  redistribute connected
  listen-range 10.0.0.0/29 peer-group pg1 limit 8
  neighbor 1.1.1.11 peer-group pg1
  neighbor 1.1.1.11 bfd
  neighbor 1.1.1.11 bfd interval 50 min-rx 5000 multiplier 50
  neighbor 1.1.1.11 update-source 10.11.12.13
  maximum-paths 8
!
address-family ipv6 unicast
  neighbor pg2 activate
  neighbor pg2 additional-paths receive
!
address-family ipv6 unicast vrf vrf11
  redistribute connected
  listen-range 8001:db8:1::/68 peer-group pg2
  neighbor 2000:db8:1:: peer-group pg2
  neighbor 2000:db8:1:: bfd
  neighbor 2000:db8:1:: bfd interval 50 min-rx 5000 multiplier 50
  neighbor 2000:db8:1:: update-source ff::ff
  neighbor 2000:db8:1:: activate
  maximum-paths 8
!
address-family l2vpn evpn
  graceful-restart
```

3. Create a BGP peer on OS ONE.

```
efa tenant service bgp peer create --name
B1 --tenant t2 --ipv4-uc-nbr 10.64.188.68,vrf2:1.1.1.11,pg1
--ipv4-uc-nbr-bfd 10.64.188.68,vrf2:1.1.1.11,true,50,5000,50 --ipv4-uc-nbr-
update-source-ip 10.64.188.68,vrf2:1.1.1.11,10.11.12.13 --ipv4-uc-dyn-nbr
10.64.188.68,vrf2:10.0.0.0/29,pg1,8 --ipv6-uc-nbr 10.64.188.68,vrf2:2000:db8:1::,pg2
--ipv6-uc-nbr-bfd 10.64.188.68,vrf2:2000:db8:1::,true,50,5000,50 --ipv6-uc-nbr-
update-source-ip 10.64.188.68,vrf2:2000:db8:1::,ff::ff --ipv6-uc-dyn-nbr
10.64.188.68,vrf2:8001:db8:1::/68,pg2

BGP Peer created successfully.
--- Time Elapsed: 4.90011281s ---
```

4. Show BGP peer on OS ONE.

```
efa tenant service bgp peer show --tenant t2
================================================================================================
================================================================================================
==============
Name        : B1
Tenant      : t2
State       : bgp-peer-created
```

```
+--------------+------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+-------------------------+-------------
+-------------+
|  Device IP   | VRF  | PeerGroup | AFI  |  SAFI   |  Remote IP   | Remote ASN |
Activate | Next Hop |   Update    |   BFD   |           BFD            |  Dev State
|  App State  |
|              |      |           |      |         |              |
|          |   Self   |  Source IP  | Enabled | [Interval,Rx,Multiplier] |
|             |
+--------------+------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+-------------------------+-------------
+-------------+
| 10.64.188.68 | vrf2 |    pg1    | ipv4 | unicast |   1.1.1.11   |            |
false   |          | 10.11.12.13 |  true   |       50, 5000, 50       | provisioned |
cfg-in-sync |
+--------------+------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+-------------------------+-------------
+-------------+
| 10.64.188.68 | vrf2 |    pg2    | ipv6 | unicast | 2000:db8:1:: |            |
false   |          |   ff::ff    |  true   |       50, 5000, 50       | provisioned |
cfg-in-sync |
+--------------+------+-----------+------+---------+--------------+------------
+----------+----------+-------------+---------+-------------------------+-------------
+-------------+
Static Peer Details

[Flag '*' indicates Multi protocol capability]


+--------------+------+------+---------+-----------------+--------+------------
+-------------+-------------+
|  Device-IP   | VRF  | AFI  |  SAFI   |  Listen Range   | Listen | Peer Group |  Dev
State  |  App State  |
|              |      |      |         |                 | Limit  |
|             |             |
+--------------+------+------+---------+-----------------+--------+------------
+-------------+-------------+
| 10.64.188.68 | vrf2 | ipv4 | unicast |   10.0.0.0/29   |   8    |    pg1     |
provisioned | cfg-in-sync |
+--------------+------+------+---------+-----------------+--------+------------
+-------------+-------------+
| 10.64.188.68 | vrf2 | ipv6 | unicast | 8001:db8:1::/68 |        |    pg2     |
provisioned | cfg-in-sync |
+--------------+------+------+---------+-----------------+--------+------------
+-------------+-------------+
Dynamic Peer Details

================================================================================
================================================================================
==============

--- Time Elapsed: 90.457906ms ---
efa tenant service bgp peer show --tenant t2 --detail
================================================================================
================================================================================
==============
Name             : B1
Tenant           : t2
State            : bgp-peer-created
Description      :

Static Peer
-----------
    Device IP                  : 10.64.188.68
```

```
    VRF                         : vrf2
    Peer Group                  : pg1
    AFI                         : ipv4
    SAFI                        : unicast
    Remote IP                   : 1.1.1.11
    Remote ASN                  :
    Update Source IP            : 10.11.12.13
    BFD Enabled                 : true
    BFD Interval                : 50
    BFD Rx                      : 5000
    BFD Multiplier              : 50
    MD5 Password                :
    Dev State                   : provisioned
    App State                   : cfg-in-sync

    Device IP                   : 10.64.188.68
    VRF                         : vrf2
    Peer Group                  : pg2
    AFI                         : ipv6
    SAFI                        : unicast
    Remote IP                   : 2000:db8:1::
    Remote ASN                  :
    Update Source IP            : ff::ff
    BFD Enabled                 : true
    BFD Interval                : 50
    BFD Rx                      : 5000
    BFD Multiplier              : 50
    MD5 Password                :
    Dev State                   : provisioned
    App State                   : cfg-in-sync


Dynamic Peer
------------
    Device IP                   : 10.64.188.68
    VRF                         : vrf2
    AFI                         : ipv4
    SAFI                        : unicast
    Listen Range                : 10.0.0.0/29
    Listen Limit                : 8
    Peer Group                  : pg1
    Dev State                   : provisioned
    App State                   : cfg-in-sync

    Device IP                   : 10.64.188.68
    VRF                         : vrf2
    AFI                         : ipv6
    SAFI                        : unicast
    Listen Range                : 8001:db8:1::/68
    Listen Limit                :
    Peer Group                  : pg2
    Dev State                   : provisioned
    App State                   : cfg-in-sync

================================================================================================
================================================================================================
==============

--- Time Elapsed: 83.928699ms ---

show running-config vrf vrf2
vrf vrf2
  member   ve 2,8192
  table-connections
```

```
    src-protocol connected dst-protocol bgp address-family ipv4
    src-protocol connected dst-protocol bgp address-family ipv6
  !
  router bgp
    local-as 4200000000
    router-id 172.31.254.98
    route-distinguisher 172.31.254.98:1
    as-notation as-plain
    address-family ipv4 unicast
      use-multiple-paths ibgp maximum-paths 8
      use-multiple-paths ebgp maximum-paths 8
      activate
    !
    address-family ipv6 unicast
      use-multiple-paths ibgp maximum-paths 8
      use-multiple-paths ebgp maximum-paths 8
      activate
    !
    instance-type evpn
      nve fab2
        l3vni 8192
        address-family ipv4 unicast
          route-target export 101:101
          route-target import 101:101
        !
        address-family ipv6 unicast
          route-target export 101:101
          route-target import 101:101
        !
      !
    !
    peer-group pg1
      remote-as 651000
      auth-password-hashed
6667b03cf1d5483bf838939d6170d6bc342281c462208ee7b56607e9353fc714
      enable-bfd profile profile_100_100_3
      update-source 10.10.10.10
      address-family ipv4 unicast
        nexthop-self
        add-paths receive
        activate
      !
      neighbor 1.1.1.11
        enable-bfd profile profile_50_5000_50
        update-source 10.11.12.13
      !
      listen-range 10.0.0.0/29 listen-limit 8
    !
    peer-group pg2
      remote-as 651000
      auth-password-hashed
6667b03cf1d5483bf838939d6170d6bc342281c462208ee7b56607e9353fc714
      enable-bfd profile profile_100_100_3
      update-source ff::ff
      address-family ipv6 unicast
        nexthop-self
        add-paths receive
        activate
      !
      neighbor 2000:db8:1::
        enable-bfd profile profile_50_5000_50
        update-source ff::ff
      !
      listen-range 8001:db8:1::/68 listen-limit 0
```

Known Limitations

Learn about the caveats for ExtremeCloud Orchestrator.

Known Limitations in Fabric Skill

Follow these caveats and limitations when using Fabric Skill.

Quality of Service (QoS) policy service support

- The XCO-driven application of policy is dynamic and can vary depending on the port's role, whether it belongs to a fabric, tenant, port channel, or tenant endpoint group.

  Tip
  As a best practice, avoid running user-driven policy operations in parallel with fabric, tenant, port channel, and tenant endpoint group operations.

  To ensure that the fabric, tenant, port channel, and tenant endpoint group configurations are effective, run the **`show`** command before proceeding with the policy operations, and vice-versa.
- Before running the force operations, including deletion, ensure that you unbind the policies (QoS) from all the relevant targets (fabric, tenant, port, port channel, and tenant endpoint group) to avoid stale policies (QoS) in the system.
- Before executing the QoS policy bind commands, remove any conflicting or additional OOB (Out of Band) QoS configurations from the switches to ensure that the correct policies are applied to the ports.
- There is no support for a lossless hardware profile. Therefore, you must switch the configuration on SLX devices to a lossy hardware profile before provisioning QoS policies from XCO.
- There is no support for egress QoS maps. While XCO allows the configuration of egress QoS maps, as a best practice, do not configure any egress QoS maps from XCO due to limitations in SLX support of egress QoS maps.

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 not 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 an **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.

Here's a sample token. Copy/Paste this in your shell:

export EFA_TOKEN=eyJhbGciOiJSUzI1NiIsImtpZCI6IjEuMCIsInR5cCI6IkpXVCJ9.eyJjb21tb25fbmFtZSI6IkVGQSBUb2tlbiBTZXJ2aWNlIiwidWFzIjpbeyJ0YXJnZXQiOiJFRkEiLCJyb2xlIjoiVlIyLVRudEFkbWluIn1dLCJvcmciOiJFeHRyZW1lIE5ldHdvcmtzIiwidmVyIjoiMS4wIiwiaWQiOiIiLCJleHAiOjE2NDUyNDcxNDIsImp0aSI6IjZjMjA4ZDUxLTkwNzgtMTFlYy1iZjk5LWNhNzk1MDY1YzIwNyIsImlhdCI6MTY0NTE2MDc0MiwiaXNzIjoiRUZBIFRva2VuIFNlcnZpY2UiLCJuYmYiOjE2NDUxNjA3NDIsInN1YiI6InVzZXIyIn0.b7m5PINijeEdNSqnTeE2ZhUrqKLKQAu079vXyBIdgHbXKt9ULfa03vMU1jfBO1qFb1-x0oHmsAQ0pSsF5JLeMaMzMflLf78ktZO8U5IePq72vM5en35IR-DNLyoGIZBeFeG6ZbBMoETzz5vf9OuefgQID3YdjcALr7yllCgDmLVFlgson77yCBpkTK15xm

```
1GRbtL7JKXZzShBE7E3kdW7N71MdM85Gc3r4l-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: Request timed out after 15 minutes. The command is still running in the background. Please verify the deletion status using 'efa tenant show' and refrain from executing any commands related to this tenant.	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, as well as any fabric operations that take more than 15 minutes, are timed out at the client side. The operation completes in the background. Run the **`efa tenant show`** command to view the actual state of the operation.

In Drift and Reconcile (DRC), if reconciliation takes over 15 minutes, the command shows `tenant-dr-timeout`. However, the tenant is still reconciling configurations in the background. As a best practice, verify using the **`efa inventory drift-reconcile execute --ip <ip_address>`** command.

Inventory device updates fail with 401 Unauthorized errors post SLX upgrade

Symptom	Condition	Workaround
After upgrading SLX, inventory device updates may intermittently fail with 401 Unauthorized errors.	The root cause is linked to a token validation issue triggered by service restarts during hardening script execution. These service restarts initiate HA events, which invalidate tokens. Manual updates succeed, but REST API calls fail due to 401 Unauthorized responses.	Log in to XCO (**`efa login`**) only after all dependent components have completed deployment and upgrade. • Do not log in to XCO (via CLI or REST) during deployment. • Avoid obtaining an `access_token` or starting an XCO session before all components are finished deploying. • Wait until the following are fully completed: ◦ XCO deployment or upgrade ◦ SLX upgrade ◦ TPVM upgrade ◦ Hardening script (HS) execution and completion