



ExtremeCloud Orchestrator GUI Administration Guide

3.2.0

9037726-00 Rev AB
January 2024



Copyright © 2024 Extreme Networks, Inc. All rights reserved.

Legal Notice

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

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

Trademarks

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

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

For additional information on Extreme Networks trademarks, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Table of Contents

Preface.....	7
Text Conventions.....	7
Documentation and Training.....	8
Help and Support.....	9
Subscribe to Product Announcements.....	9
Send Feedback.....	10
What's New in this Document.....	11
Welcome to ExtremeCloud Orchestrator.....	13
Fabric Automation and Orchestration (Fabric Skill).....	13
Visibility Solution (Visibility Skill).....	14
Packet Broker Functions.....	15
XCO Limitations.....	16
XCO Deployment.....	17
Navigate the User Interface.....	18
Log in to XCO.....	18
User Interface.....	19
Refresh Page View.....	20
Pagination.....	21
Search, Group, and Filter.....	21
Dashboard.....	22
System Widget.....	22
Fabrics Widget (Fabric Mode).....	23
Locations Widget.....	23
Devices Widget.....	24
Users Widget.....	24
Help & Support Widget.....	24
Support Save.....	25
Register Remote Server.....	25
Generate Support Save.....	25
Download Support Save.....	26
Locations.....	27
Add Location.....	27
Download Location Definition File.....	29
Delete Location.....	29
Display Location-Specific Device List.....	29
Display Locations Map View.....	29
Device Inventory.....	31
Device Credentials.....	31

Add Devices	31
Create a Device Definition File.....	33
Download Bulk Device Inventory.....	33
Delete Device.....	34
Overview (Packet Broker Mode).....	34
Device Statistics.....	34
View Statistics in a Device Dashboard.....	35
Device Actions (Packet Broker Mode).....	37
Save the Running Configuration of SLX and MLX Devices.....	37
Refresh Configuration.....	37
Export Configuration.....	38
Packet Capture.....	39
Clear Counters.....	42
View Logs.....	43
Delete a Device from the Device Overview Page.....	44
Policies and Configuration (Packet Broker Mode).....	44
Policies.....	44
Policy Rule Matches.....	48
UDA Profiles.....	52
Ingress Groups.....	55
Egress-Group.....	58
Egress.....	60
Mirrors.....	61
Tunnels.....	62
Quality of Service.....	66
Port Channels.....	67
Ports.....	70
Truncation Profile.....	71
TACACS+ Authentication.....	72
Slots.....	73
Parts Details (Packet Broker Mode).....	75
SLX Optical Statistics.....	75
View Device Inventory.....	75
Download Device Inventory.....	76
Monitor/Troubleshoot (Packet Broker Mode).....	76
Troubleshoot Configuration.....	77
Firmware Upgrade.....	79
Register Firmware Host.....	79
View Registered Firmware Hosts.....	80
Change a Firmware Host.....	81
Delete a Firmware Host.....	81
Upgrade Firmware (Device Level)	81
Rollback Upgrade (Packet Broker Mode).....	83
Users.....	84
Role Based Access Control.....	84
User Roles.....	85
Authentication Tokens.....	85
Local.....	86
Add User.....	86

Edit User.....	86
Block User.....	87
Unblock User.....	87
Request Reset Password.....	87
Change Password on First Login.....	88
Delete User.....	88
Host.....	88
Change Host User Role.....	88
User Settings.....	89
Authentication Settings.....	89
LDAP Settings.....	90
TACACS+ Settings.....	92
Change a Server Configuration.....	94
Delete a Server Configuration.....	94
Change Password.....	94
Logout.....	95
Logs.....	96
System Logs.....	96
User Logs.....	97
Library (Packet Broker Mode).....	99
Matches.....	99
Create a Policy Rule Match in the Library.....	99
Change a Policy Rule Match in the Library.....	102
Export a Policy Rule Match from the Library.....	102
Clone a Policy Rule Match.....	103
Delete a Policy Rule Match from the Library.....	103
Policies.....	103
Create a Policy in the Library.....	103
Change a Policy in the Library.....	104
Export a Policy.....	104
Clone a Policy.....	104
Delete a Policy in the Library.....	105
UDA.....	105
Create an MLX UDA Profile in the Library.....	106
Create an SLX UDA Profile in the Library.....	106
Change a UDA Profile in the Library.....	107
Delete a UDA Profile in the Library.....	107
Export a UDA Profile.....	107
Fabrics (Fabric Mode).....	108
Create a Non-Clos Fabric.....	109
Create a 3 Stage Clos Fabric.....	113
Create a 5 Stage Fabric.....	118
View Fabric Topology.....	124
Download Fabric Inventory.....	125
Delete Fabric.....	126
Download Health Report.....	126
Fabric Actions.....	126
Edit Fabric Topology.....	127

Delete a Fabric from the Fabric Actions Menu.....	128
Network Essentials.....	129
Firmware Upgrade.....	130
Clone a Fabric.....	135
FAQs.....	137
Where are Visibility skill logs located?.....	137
Where are the Inventory Service logs located?.....	137
Where are the Installer logs located?.....	137
What are some common reasons for installation failures?.....	137
Why does the web user interface not load on the browser?.....	137
What are some common reasons for XCO log-in failures?.....	137
Where are authentication failures captured?.....	138
What are possible reasons for device registration failures?.....	138
Why is there a delay in loading the dashboard or statistics in the web UI?.....	138
Why is the device configuration blocked from the web UI?.....	138
What are possible reasons for configuration failures?.....	138
How do I check that all services are up and running?.....	138
Why are the device syslogs not visible?.....	139
How to collect the SupportSave data for troubleshooting?.....	139



Preface

Read the following topics to learn about:

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

Text Conventions

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

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

Table 1: Notes and warnings






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

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
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.
<i>italic text</i>	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, <i>member [member...]</i> .
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

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

[Extreme Optics Compatibility](#)

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

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

Help and Support

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

Extreme Portal

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

The Hub

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

Call GTAC

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

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

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

Subscribe to Product Announcements

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

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.

3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

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

Send Feedback

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

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, do either of the following:

- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

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



What's New in this Document

The following table describes changes to this guide for the ExtremeCloud Orchestrator 3.2.0 release.

Table 4: Summary of changes

Feature	Description	Link
Common landing page	The dashboard or the common landing page provides an overview of system health.	<ul style="list-style-type: none">• Dashboard on page 22• User Interface on page 19
Fabric management	ExtremeCloud Orchestrator automates and orchestrates SLX IP fabrics through CLI and user interface.	<ul style="list-style-type: none">• Dashboard on page 22• Fabrics (Fabric Mode) on page 108
Network essential configurations	ExtremeCloud Orchestrator supports network essential configurations that are required for creating and configuring fabric networks.	<ul style="list-style-type: none">• Network Essentials on page 129
Support Save	You can generate and download Support Save logs for troubleshooting.	<ul style="list-style-type: none">• Support Save on page 25
Bulk device inventory download	You can download multiple or bulk device inventory information.	<ul style="list-style-type: none">• Download Bulk Device Inventory on page 33
SLX optical statistics	ExtremeCloud Orchestrator collects the device inventory and channel media information during device discovery and stores it in the database.	<ul style="list-style-type: none">• SLX Optical Statistics on page 75

Table 4: Summary of changes (continued)

Feature	Description	Link
User management enhancements	ExtremeCloud Orchestrator supports the following user management enhancements: <ul style="list-style-type: none"> • Multiple user role mapping • Host user role change 	<ul style="list-style-type: none"> • Add User on page 86 • Change Host User Role on page 88
Role Based Access Control (RBAC)	Role Based Access Control (RBAC) defines the capabilities that a user account has based on the assigned role.	<ul style="list-style-type: none"> • Role Based Access Control on page 84 • Map an LDAP User Role on page 92 • Map a TACACS+ User Role on page 93

For more information about this release, see the [ExtremeCloud Orchestrator Release Notes, 3.2.0](#).



Welcome to ExtremeCloud Orchestrator

[Fabric Automation and Orchestration \(Fabric Skill\)](#) on page 13

[Visibility Solution \(Visibility Skill\)](#) on page 14

[XCO Limitations](#) on page 16

[XCO Deployment](#) on page 17

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

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

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

For information about evolution of EFA and XVM into XCO, see the [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).



Note

All procedures in this document are performed through GUI.

Fabric Automation and Orchestration (Fabric Skill)

XCO automates and orchestrates SLX IP fabric networks through CLI and UI.

For more information about fabric skill, see:

- [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#)
- [ExtremeCloud Orchestrator Command Reference, 3.2.0](#)
- [ExtremeCloud Orchestrator Deployment Guide, 3.2.0](#)
- [ExtremeCloud Orchestrator Security Guide, 3.2.0](#)

- [ExtremeCloud Orchestrator Hyper-V Integration Guide, 3.2.0](#)
- [ExtremeCloud Orchestrator VMware vCenter Integration Guide, 3.2.0](#)

Visibility Solution (Visibility Skill)

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

Although devices have different functionality and different configuration methods, XCO seamlessly interacts with all supported devices for simplified management.

XCO managed objects work together to accomplish packet broker functions. You can configure these objects from the user interface. For more information, see [Packet Broker Functions](#) on page 15.

Table 5: Managed objects

Object	Description
Ports and port channels	The interfaces on which traffic enters and exits a device. You can associate ports and port channels with ingress groups and egress. For more information, see Ports on page 70 and Port Channels on page 67.
Egress	A port or port channel that you associate with an egress policy, which identifies the actions to take on egress traffic. For more information, see Create an Egress for Devices on page 60.
Egress group	A set of interfaces and ports on which traffic is forwarded after a policy is applied. For more information, see Create an Egress Group on page 58.
Ingress group	A collection of ports, port channels, and tunnels on which monitored traffic is received. You can select several actions to perform on the incoming traffic and you can associate the ingress group with an ingress policy. For more information, see Ingress Groups on page 55.
Policy rule matches	The parts of a packet header that a rule targets, such as the source port or the payload length. One or more rules constitute a match. You associate matches with ingress or egress policies. For more information, see Policy Rule Matches on page 48.

Table 5: Managed objects (continued)

Object	Description
Ingress policy (or route map)	The actions to apply to inbound packets. You can associate policy rule matches and egress groups, and select other actions such as packet slicing and scope shift. For more information, see Create an Ingress Policy for a Device on page 46.
Egress policy (or listener policy)	The actions to apply to outbound packets. You can associate policy rule matches and select other actions such as packet slicing and header stripping. For more information, see Create an Egress Policy for a Device on page 44. Note: Applies to 9920 devices only.
User-defined access list (UDA)	The UDA profiles for SLX and MLX devices. For more information, see UDA Profiles on page 52.
Transport tunnel termination and encapsulation	The GRE or ERSPAN tunnels to associate with ingress groups or egress. For more information, see Tunnels on page 62. Note: Applies to 9920 devices only.
Quality of Service (QoS)	The QoS configuration is used on 9920 devices to manage traffic delivery. For more information, see Quality of Service on page 66.
Mirrors	A copy of the egress port traffic on 9920 is forwarded to the configured mirror destination port. For more information, see Mirrors on page 61.

Packet Broker Functions

A network packet broker aggregates network traffic from multiple ports for forwarding to analysis applications.

When a packet broker is attached to networking devices, a copy of the traffic that passes through the devices is sent to the packet broker. Based on your configuration, the packet broker filters the copied traffic for the data that you want to analyze. The packet broker then sends the filtered traffic to an analysis application.

In general, packet brokers can perform the following types of actions on copied network traffic.

Table 6: Packet broker functions

Function	Description
ACL filtering	Directs network traffic based on Layer 2 to Layer 4 protocol headers
Aggregation	Combines traffic that from multiple ports and directs it to one port or port channel
Decapsulation	Removes the outer tunnel headers from a packet
Header stripping	Removes header tags that are not supported by some visibility applications, including 802.1BR, VN (virtual NIC), VLAN, VXLAN, GTPU, GRE, and IPIP headers
Load balancing	Distributes network traffic among ports in a port channel
Packet slicing	Truncates packets to a specific size across ports.
Replication	Copies network traffic to multiple ports and port channels.
Route map forwarding	Redirects Layer 2 and Layer 3 packets to the selected physical or port channel interface
Transport tunnel termination	<ul style="list-style-type: none"> • GRE (Generic Routing Encapsulation). Creates a tunnel that encapsulates (or wraps) packets that use one type of protocol inside packets that use a different protocol. • ERSPAN (Encapsulated Remote Switched Port Analyzer): Creates a tunnel that mirrors traffic from source ports for delivery to destination ports on a different device.
Transport tunnel encapsulation	GRE only

XCO Limitations

XCO limitations are as follows:

- Hostname or DNS name based device discovery is not supported.
- Device location cannot be modified after discovery.
- CLI support is not available for visibility skill.
- Only live statistics data streaming is supported.
- Secured Syslog configuration is not supported for MLX devices.

- User-defined access list (UDA) configuration is not migrated while configuration migration from an SLX or MLX to 9920.
- Listener policy byte count is incorrect for 9920 when truncation is enabled.
- Special characters such as %, { }, \, and = are not supported in Name fields.
- If a device configured with both IPv4 and IPv6 addresses is discovered, only one entry is added to XCO. The first discovered IP address is used for communicating with that device.
- Device discovery failure is not listed on the Device log page for non-packet broker devices in the packet broker mode.
- All configurations are reverted when a port channel deployment fails. However, a LAG is created and deleted immediately, and the events are captured in the device logs.
- Firmware upgrade requires an absolute path to the image location.

XCO Deployment

For information about deploying XCO, see the [ExtremeCloud Orchestrator Deployment Guide, 3.2.0](#).



Navigate the User Interface

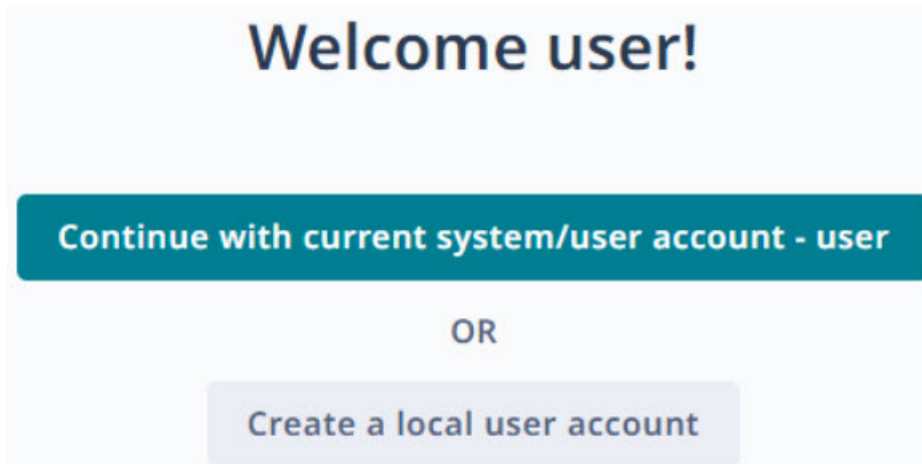
- [Log in to XCO](#) on page 18
- [User Interface](#) on page 19
- [Refresh Page View](#) on page 20
- [Pagination](#) on page 21
- [Search, Group, and Filter](#) on page 21

You can access XCO using the latest two versions of Google Chrome or Microsoft Edge web browsers.

Log in to XCO

Procedure

1. In a web browser, open `http://xx.xx.xx.xx/login`, where `xx.xx.xx.xx` is the IP address of the control plane node.
2. Complete the **Username** and **Password** fields.
3. Select **Login**.



If this is your first login as a host user, you are prompted to either continue with the current host user account or create a new local user account. Otherwise, the user interface opens to the **Dashboard** page.

Follow the instructions in [Add User](#) on page 86 to create new user accounts.

Local users are prompted to reset the password on first login.

Related Topics

[Add Location](#) on page 27

The Location Definition file (in CSV format) identifies geographical locations.

[Add Devices](#) on page 31

User Interface

The XCO interface provides access to all system functions.

[Table 7](#) describes the numbered elements in this diagram.

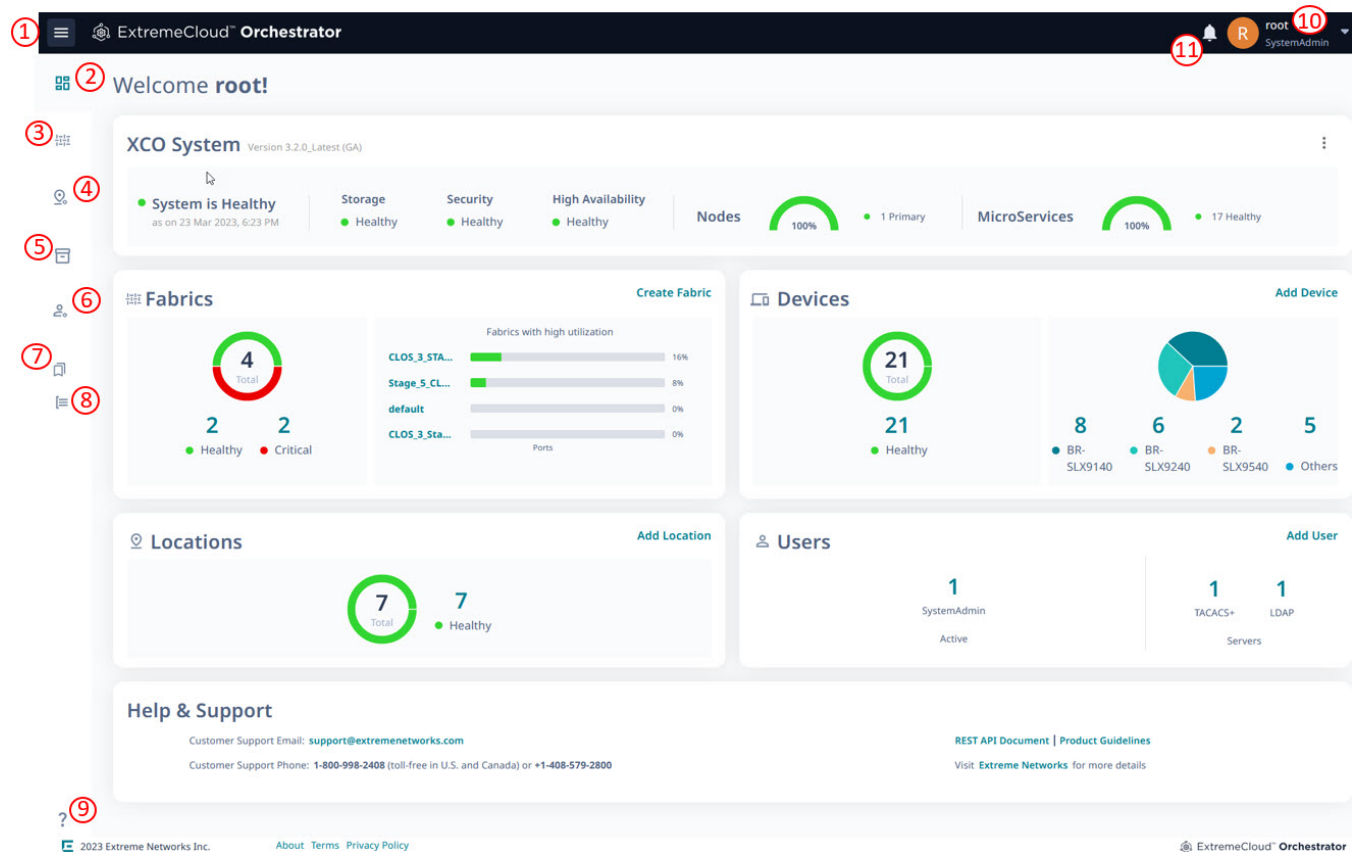



Figure 1: XCO user interface


Table 7: User interface descriptions

Legend	Interface Area	Description
1	Navigation menu	Provides access to all pages of the interface.
2	Dashboard	Provides an overview of system health and quick access to locations, devices, fabrics, and users configuration pages. See Dashboard on page 22.

Table 7: User interface descriptions (continued)

Legend	Interface Area	Description
3	Fabrics (Fabric Mode only)	Provides access to the fabrics management page. See Fabrics (Fabric Mode) on page 108.
4	Locations	Provides access to the location management page. See Locations on page 27.
5	Device Inventory	Provides physical details and access to all configuration settings of the selected device. Details vary by device type. See Device Inventory on page 31.
6	Users	Provides access to settings for users, profile, and authentication. See Users on page 84.
7	Logs	Provides access to the logs page. See Logs on page 96.
8	Library (Packet Broker Mode only)	Provides access to configured matches, policies, and UDA (user-defined ACL) profiles. See Library (Packet Broker Mode) on page 99.
9	Help & Support	Provides access to the help and support information.
10	User Profile	Displays the username and role of the logged-in user. From here, you can perform the following tasks: <ul style="list-style-type: none"> • Change own password • Log out
11	Notifications	Provides access to the notifications page. The notifications are user specific and do not persist. The  icon indicates new notifications.

Refresh Page View


When you add a new entry or modify an existing entry in a table in the XCO user interface, you are prompted to refresh () the page to view the latest changes.

Pagination

XCO supports pagination in all pages that show detailed data, such as locally configured users, devices, device configurations, policies, authentication servers, and locations.

Procedure

Select the required **Page Size (5, 10, 20, 50, 100)** to specify the number of entries in a table.

- The default page size is 10.
- Use the **Previous** and **Next** icons () to scroll through the list.

Limitation:

The user interface displays incorrect data on the previous page when you scroll through list pages after applying filters.

Search, Group, and Filter



You can search for an item and organize lists in the XCO user interface.


About This Task

You can group records based on the pre-defined criteria that vary for different windows.

Use the **Previous** and **Next** icons () to scroll through the lists.

Procedure

1. To search for a record in a page, enter a search attribute such as object name, IP address, location in the **Search** field and click **Search** ().
To clear the search, click **X** in the **Search** field.
2. To group records in a page, select **Group By** and choose an attribute.
To clear the grouping, select the **Clear** option.
The list is organized by the grouping attribute you selected. The headings are collapsible.
3. To filter records in a page, select **Filter** () and choose the filter attribute.
To clear an individual filter, click **x** for the appropriate filter. To clear all the filters, click **Clear All Filters**.
The list is organized by the filtering attribute you selected.



Dashboard

- [System Widget](#) on page 22
- [Fabrics Widget \(Fabric Mode\)](#) on page 23
- [Locations Widget](#) on page 23
- [Devices Widget](#) on page 24
- [Users Widget](#) on page 24
- [Help & Support Widget](#) on page 24
- [Support Save](#) on page 25

The XCO's **Welcome user!** dashboard screen or the landing page provides an overview of system health and provides quick access to various pages such as Fabrics, Locations, Devices, and Users. The critical errors in the system are marked in red.

The dashboard varies depending on the logged-in user role. For more information about user roles, see [User Roles](#) on page 85.

System Widget

The system widget on the dashboard displays information about nodes and microservices running in the system, health status of storage, security, and high availability. It also provides access to the **Support Save** menu.




Table 8: System widget components

Component	Description
Storage	Indicates the storage status.
Security	Indicates the security status.
High Availability	Indicates the high availability status.
Nodes	Indicates the count of primary and standby nodes.
Microservices	Indicates the count of healthy, degraded, and critical state of microservices.

Related Topics

[Support Save](#) on page 25

XCO supports Support Save logs collection for troubleshooting.

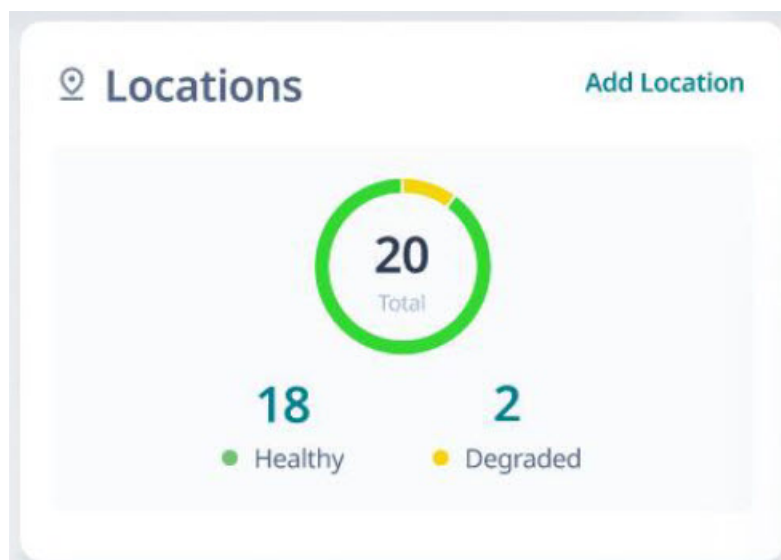
Fabrics Widget (Fabric Mode)

The **Fabrics** widget on the dashboard displays an overview of fabrics health and the five most heavily used fabrics with high utilization. Use the Fabrics widget to access the Fabrics management page.



Locations Widget

The **Locations** widget on the dashboard displays the total number of locations and their health status. Use the locations widget to access the Locations management page.



Devices Widget

The **Devices** widget on the dashboard displays the total number of discovered devices and their health status along with type specific device health status. Use the devices widget to access the `Devices` management page.



Users Widget

The **Users** widget on the dashboard displays the number of active users, active users by type, TACACS+ servers, and LDAP servers information. Use the users widget to access the `Users` management page.



Help & Support Widget

The **Help & Support** widget displays customer support contact information.

Support Save

XCO supports Support Save logs collection for troubleshooting.

You can generate and download Support Save logs as follows:

1. Generate system Support Save logs
2. Configure remote server for copying Support Save logs
3. Download Support Save logs

Related Topics

[Register Remote Server](#) on page 25

[Generate Support Save](#) on page 25


[Download Support Save](#) on page 26

Register Remote Server

About This Task

You can configure a remoter server to copy the generated Support Save logs.

Procedure

1. In the Navigation menu, select **Dashboard**.
2. Select  in the System Health widget.
3. Select **Support Save**.
4. Select **Register**.
5. In the **IP Address** field, add the IP address of the remote server.
6. In the **Username** and **Password** fields, add the device credentials.
7. In the **Directory** field, provide the remote server path.
8. In the **Protocol** field, select the protocol.
 - **FTP**
 - **SCP**
9. In the **Max Support File Save Limit**, select a value to configure the number of support save files.


When the configured support save file limit is reached, the oldest support save file is deleted when a new support save request is triggered.

- The number of save files defaults to five and a maximum of 20 files are supported.
- A minimum of two support files are required.

10. Select **Register**.

Generate Support Save

Procedure

1. In the Navigation menu, select **Dashboard**.
2. Select  in the System Health widget.

3. Select **Support Save**.
4. Select **Generate Support Save**.

The new support save file is added to the list of support save files.

Support Save Files

Generate Support Save

Showing 1 - 4 of 4 results Page Size Latest as on 12:17:58 PM

Name	Time	Status	Actions
efa_2023-03-02T12-21-29.336.logs.zip	2023-03-02	● Completed	
efa_2023-03-02T12-21-01.625.logs.zip	2023-03-02	● Completed	
efa_2023-03-02T12-20-55.129.logs.zip	2023-03-02	● Completed	
efa_2023-03-02T11-32-49.508.logs.zip	2023-03-02	● Completed	

< 1 >

When the configured support save file limit is reached, the oldest support save file is deleted when a new support save request is triggered.

- The number of save files defaults to five and a maximum of 20 files are supported.
- A minimum of two support files are required.

For information on configuring the support save file limit, see [Register Remote Server](#) on page 25.

5. Select **OK**.

A notification is displayed when the Support Save file is generated.

Download Support Save

Procedure

1. In the Navigation menu, select **Dashboard**.
2. Select in the System Health widget.
3. Select **Support Save**.
4. Select **Download** () for the required support save logs file.

The selected support save file is downloaded to your device.



Locations

[Add Location](#) on page 27

[Download Location Definition File](#) on page 29

[Delete Location](#) on page 29

[Display Location-Specific Device List](#) on page 29

[Display Locations Map View](#) on page 29

XCO **Location Management** allows you to view and manage devices from different geographical locations. A default location is created during the XCO boot up which can be used for small deployments.

XCO manages the region, site, or location information for categorizing the devices by their physical location.



Note

- The default location cannot be modified or deleted.
- When an existing location is deleted, all its devices are moved to the default location.
- The device location cannot be modified after discovery.

Add Location

The Location Definition file (in CSV format) identifies geographical locations.

About This Task

After XCO is installed, you can upload the CSV file to the interface. For information about deploying XCO, see the [ExtremeCloud Orchestrator Deployment Guide, 3.2.0](#).

Procedure

1. In the Navigation menu, select **Locations**.

Showing 1 - 7 of 7 results

Search Group By None Page Size 10 Latest as on 6:49:12 PM Refresh Download Filter

Name	Address	City	Zipcode	Country	Type	Region	Actions
Toronto, Ontario	Toronto, Ontario	Toronto, Ontario	03079	US	Engineering,Testing,Lab	North America	Edit Delete
Salem, NH	Salem, NH	Salem, NH	03079	US	Engineering,Testing,Lab,Camp	North America	Edit Delete
San Jose	San Jose	San Jose	95119	US	Data center,Campus,Lab,Engin	North America	Edit Delete
Mumbai	Mumbai	Mumbai	400099	IN	Data center	Asia	Edit Delete
Bengaluru	Bengaluru	Bengaluru	560068	IN	Data center,Campus,Lab,Engin	Asia	Edit Delete
Pune	Pune	Pune	411028	IN	Engineering,Testing	Asia	Edit Delete
default	-	-	-	-	-	-	Edit

< 1 >

The **Location Management** window opens.

2. Select **Add Location**.

The **Add New Location** window opens.

3. To add new locations manually, take the following steps:
 - a. Select **Add Address** and type the following information:
 - Name
 - Type
 - Region
 - Street Address
 - Country
 - State
 - City
 - Zipcode
 - Latitude
 - Longitude



Note

All the above mentioned fields are mandatory to add a new location.


- b. Select **Save**.
4. To import the `locations.csv` file, do the following:
 - a. Select **Import Location**.
 - b. Click **Select File**.

Use the sample .CSV file provided to create a .CSV file with all the location details.
 - c. Upload the .CSV file.
 - d. Select **Save**.

Download Location Definition File

The Location Definition file (in CSV format) identifies regions and their associated zones and managed locations.

Procedure

1. In the Navigation menu, select **Locations**.
2. Select  **Download**.



A file in .csv format is downloaded to your device.

Delete Location

About This Task


When an existing location is deleted, all associated devices are updated and moved to the default location.

Procedure

1. In the Navigation menu, select **Locations**.
2. In the **Location Management** page, select **Delete** () from the Actions column () for the location you want to delete.


Display Location-Specific Device List

Procedure

1. In the Navigation menu, select **Locations**.
2. In the **Location Management** page, click anywhere in the location row except the Actions column () to display the list of devices associated with the location.
To configure and manage devices, see [Device Inventory](#) on page 31.

Display Locations Map View

Procedure

1. In the Navigation menu, select **Locations**.
2. In the upper right corner of the **Location Management** page, select  to display the map view.

The default map view is the list view.

← Location Management

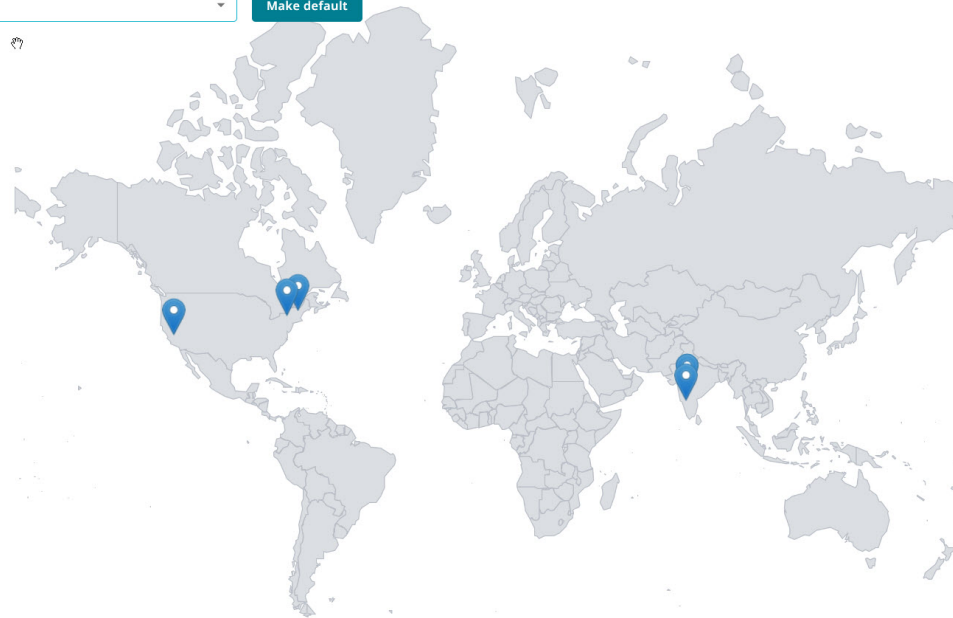
+ Add Location



Select Country (Optional)

World

Make default



3. Select a country from the drop-down menu to view the country specific location information.
4. (Optional) Select **Make Default** to make the selected country view as the default map view.



Device Inventory

- [Device Credentials](#) on page 31
- [Add Devices](#) on page 31
- [Create a Device Definition File](#) on page 33
- [Download Bulk Device Inventory](#) on page 33
- [Delete Device](#) on page 34
- [Overview \(Packet Broker Mode\)](#) on page 34
- [Device Actions \(Packet Broker Mode\)](#) on page 37
- [Policies and Configuration \(Packet Broker Mode\)](#) on page 44
- [Parts Details \(Packet Broker Mode\)](#) on page 75
- [Monitor/Troubleshoot \(Packet Broker Mode\)](#) on page 76
- [Firmware Upgrade](#) on page 79

ExtremeCloud Orchestrator supports device discovery based on IP address, user credentials, and location information.

Device discovery limitations are as follows:

- Hostname or DNS name based device discovery is not supported.
- Device location cannot be modified after discovery.
- If a device configured with both IPv4 and IPv6 addresses is discovered, only one entry is added to ExtremeCloud Orchestrator. The first discovered IP address is used for communicating with that device.

Device Credentials

The device credentials are stored in the Inventory Service database. All other microservices retrieve device credentials from the Inventory Service.

Add Devices

Before You Begin

- To be able to add multiple devices in bulk, create a Device Definition File, a CSV file that specifies the devices that you want to add. For more information, see [Create a Device Definition File](#) on page 33.
- The MLX devices must be configured for SSH as they are not AAA enabled and do not have the default user name and password.

About This Task

When a device is discovered, the device state is updated as `In Progress`. If the device connection is not successful, the appropriate error message is added to the notifications page.



Note

XCO deployed in packet broker mode supports device discovery notifications only for packet broker devices.

Procedure

1. In the Navigation menu, select **Device Inventory > Add Devices**.
2. Proceed to step 3 to add devices manually. Else, go to step 4 on page 32 to add multiple devices in bulk.
3. Select **Manually** and complete the following fields to add devices manually:
 - a. In the **Add List of IP(s)** field, enter the IPv4 or IPv6 address of the devices.
You can add a single IP address or a list of IP addresses enclosed in double quotes as shown in the following examples:

```
1.1.1.1  
"1.1.1.1, 2.2.2.2"
```
 - b. In the **Location** field, select the location where the device resides.
 - ExtremeCloud Orchestrator 3.2.0 deployed in IP fabric mode supports only the **default** location.
 - XCO creates periodic system backup at scheduled intervals and all services are locked during system backup. For more information, see the [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).

The location drop-down list will not be available during system backup. This is reflected in the user interface as "Service is Locked with reason backup".
 - c. Enter the **Username** and **Password** information.
 - d. (9920 only) In the **LACP System Priority** field, select a value to set the LACP system priority.
4. Select **Import > Select File** to browse to the CSV file.
A sample CSV file template is available for download to create device definition files.
5. Select **Save**.

Create a Device Definition File

A Device Definition file (in CSV format) identifies devices by data such as IP address, location, and credentials.

About This Task

You use a Device Definition file to add multiple devices in bulk. Each row in the CSV file has a variation of the following format.

```
IP_ADDRESS, USER_NAME, PASSWORD, LOCATION, LACP_SYSTEM_PRIORITY
```

Table 9: Field descriptions

Field Number	Field	Description
1	IP_ADDRESS	One or more IPv4 or IPv6 addresses, separated by commas.
2	USER_NAME	Credentials for accessing the device, and not necessarily the credentials of the default user.
3	PASSWORD	Credentials for accessing the device, and not necessarily the credentials of the default user.
4	LOCATION	Specifies the name of a location.
5	LACP_SYSTEM_PRIORITY	Specifies LACP system priority (9920 only) .

Procedure

1. Create a CSV file with a file name of your choosing.
Use the **Sample CSV** file available at **Device Inventory > + Add Devices > Import** to create the .CSV file.
2. Add content to the .CSV file.
3. Save the CSV file to a location that is accessible from the XCO user interface.

Example



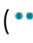
```
IP_ADDRESS,USER_NAME,PASSWORD,LOCATION,LACP_SYSTEM_PRIORITY
2620:100:c:fe08::110,admin,password,Site1,
2620:100:c:fe08::111,admin,password,Site1,
10.37.128.70,admin,password,Site1,
```

Download Bulk Device Inventory

You can download multiple or bulk device inventory information.


Procedure

1. In the Navigation menu, select **Device Inventory**.

2. In the **Devices** page, select **Download** ()
A zip file containing individual CSV files for each device type is downloaded.
3. To download the inventory of selected devices, do the following:
 - a. Select the check boxes for the devices you want to download.
 - b. Select **Download** ().
 - Alternatively, you can select **Download Inventory** from the Actions column () for the required device.
 - A zip file containing individual CSV files for each device type is downloaded.

Delete Device

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the Devices page, select **Delete Device** from the Actions column () for the device you want to delete.
3. Select **Confirm** when prompted.

Overview (Packet Broker Mode)

You can view real-time device and health statistics in the device-specific dashboard.

Device Statistics

Device Statistics

The statistics are obtained from supported devices in the following ways.

- **Extreme 9920 devices:** When a 9920 device is discovered, XCO uses gNMI to subscribe to the required statistic types. The collected statistics are processed and displayed in dashboards.
- **SLX devices:** When an SLX device is discovered, XCO is configured as a telemetry collector for streaming statistics. Streamed statistics are processed and displayed in dashboards.
- **MLX devices:** XCO runs CLI commands periodically to collect statistics, which are processed and displayed in dashboards.

Table 10: Statistics collection interval by device type

Statistic	Interval (seconds)		
	9920	SLX	MLX
System	10	60	180
Interface	10	10	180
Interface summary	10	10	180

Table 10: Statistics collection interval by device type (continued)

Statistic	Interval (seconds)		
	9920	SLX	MLX
Ingress policy	10	60	180
Egress policy	10	NA	NA
Ingress group	10	NA	NA
Egress group	10	NA	NA
Transport tunnel	10	NA	NA
Tunnel encap	10	NA	NA

Table 11: Supported statistics by device type

Statistic	9920	SLX	MLX
System	Yes	Yes	Yes
Interface	Yes	Yes	Yes
Interface summary	Yes	Yes	Yes
Ingress policy	Yes	Yes	Yes
Egress policy	Yes	No	No
Ingress group	Yes	No	No
Egress group	Yes	No	No
Transport tunnel	Yes	No	No
Tunnel encap	Yes	No	No

View Statistics in a Device Dashboard

The reports on the device dashboard provide real-time, per-device statistics.

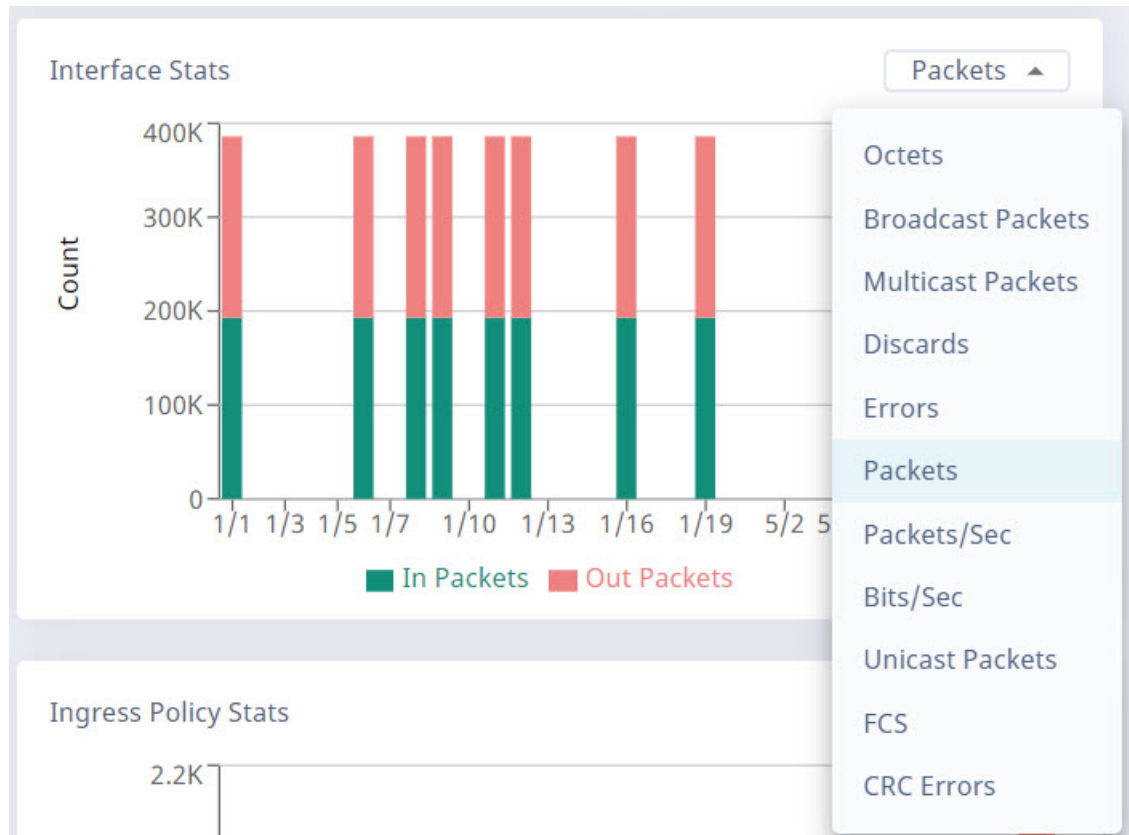
Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.

The **Overview** tab displays several device-specific default reports.

- To view a different statistic in a report, select the statistic from the list in the upper right corner of the report.

Figure 2: Statistics list



- To view statistics details, hover your cursor over an item in a report.

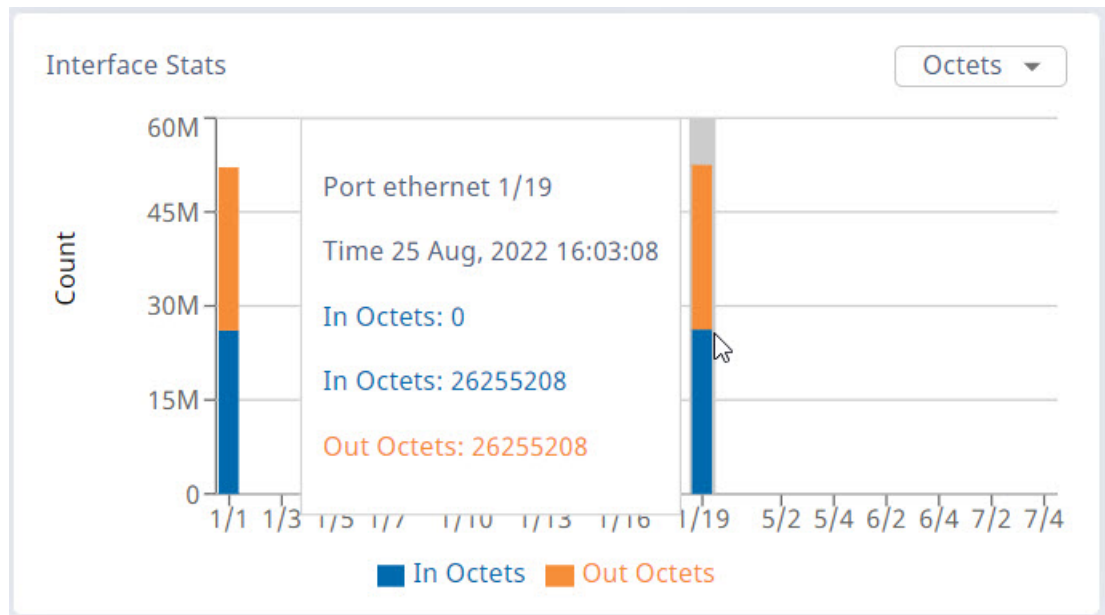


Figure 3: Statistics details

Device Actions (Packet Broker Mode)

You can perform the following tasks from the **Device Actions** menu on the Device **Overview** page.

- Save Running Configuration
- Refresh Configuration
- Export Configuration
- Capture Packets
- Clear Counters
- View Logs
- Upgrade Firmware
- Delete

Save the Running Configuration of SLX and MLX Devices

You can save the running configuration of SLX and MLX devices as start-up configuration for devices.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select **Save Running Configuration** from the Device Actions menu.

A message is displayed indicating saving configuration. The running configuration is saved to the `startup config` file.

Refresh Configuration

You can use the refresh function to retrieve the latest configuration from a device. If there are any failures, the Notification page is updated.

About This Task

Perform this procedure to sync the configuration of a device with XCO.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.

3. In the Device Actions menu, select **Refresh Configuration**.

**Note**

The XCO user interface does not support hostname updates. If you update device hostnames directly on the device, remove and re-add those devices in the user interface to view the updates.

A message is displayed that the device is in the **in_progress** mode and to wait until the device becomes **healthy**.

Export Configuration

You can export the configuration of an SLX or MLX device to an Extreme 9920 device.

About This Task

In general, the following configuration is exported:

- Policy rule matches (IPv4, IPv6, and L2 only)
- Policies
- Ingress group
- Egress
- Egress group

The following configuration items are not exported. These items appear in red text on the Export Configuration page.

- Special characters such as %, {, }, \, and = are not supported on the Extreme 9920 device. Policies and rule matches are not exported if the names of those items contain these special characters.
- User-defined access lists (UDA) are not supported on the Extreme 9920 device and are not exported.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Actions menu, select **Export Configuration**.
4. In the **To Device** field, select the device to which you want to export the configuration.
 - The configuration of the source device is displayed.
 - Items in red text under each Configuration drop-down menu are not exported.
 - Items flagged with an "i" symbol require the selection of one or more ports before you can export the items.
5. Select all required ports.
6. Select **Save**.

The configuration is exported to the destination device.

Packet Capture

XCO supports Packet Capture (PCAP) on 9920 and SLX devices.

PCAP captures packet data from the traffic that enters and leaves a device. The captured packets are stored in one or more PCAP files. After capturing the configured number of packets, packet capturing process automatically stops for the selected interface.

Ingress packets are captured before processing and egress packets are captured after processing, including header alterations.

You can use the data in PCAP files to monitor and analyze network traffic for information such as bandwidth usage, DNS resolution, network intrusion, and debugging.

- The packets received from data-path are written to the active PCAP file, `pktcapture_running.pcapng`.
- The active PCAP file is renamed and saved as `pktcapture_N.pcapng`, where N is 1-25.
- A maximum of 25 PCAP files with a file size of 100 MB each is supported for 9920. Packet capture automatically stops when 25 PCAP files are available. The existing PCAP files have to be removed to restart packet capture.
- The capture writes to the active PCAP file until file size reaches 100 MB. The PCAP file is then renamed and saved.
- Every SLX packet capture overwrites the previous PCAP file.
- If the capture is manually stopped, irrespective of the current file size, the active PCAP file is renamed and saved.
- (9920 only) XCO supports 10 simultaneous packet captures.

Start a PCAP on SLX Devices

PCAP information from an SLX device is captured in a file that you can download.

Before You Begin

Because PCAP for SLX devices is supported on only one port at a time, you must stop an existing PCAP before you can begin a new one. For more information, see [Stop a PCAP](#) on page 41.

About This Task

Every SLX packet capture overwrites the previous PCAP file.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select **Capture Packets** from the Device Actions menu.
4. Select **Configure Port Capture** or **+ (Enable PCAP)** as appropriate to start a new packet capture.
5. From the Packet Capture actions, select **Start**.

6. In the **Port** field, select a port on which to capture packets.
7. In the **Direction** field, select the packet type:
 - **Ingress (RX)**
 - **Egress (TX)**
8. In the **Packet Count** field, select the number of packets that you want to capture, from 1 to 8,000.

Packet capture stops when the selected number of packets has been captured.
9. Select **Add**.

The capture configuration is displayed on the right.
10. Select **Save**.

The PCAP file is added to the Packet Capture page.

Start a PCAP on Extreme 9920 Devices

PCAP information from a 9920 device is displayed in the XCO interface.

About This Task

A maximum of 25 PCAP files are supported for 9920.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select **Capture Packets** from the Device Actions menu.
4. Select **Configure Port Capture** or **+ (Enable PCAP)** as appropriate to start a new packet capture.
5. From the Packet Capture actions, select **Start**.
6. In the **Port** field, select a port on which to capture packets.
7. In the **Direction** field, select the packet type:
 - **Ingress (RX)**
 - **Egress (TX)**
 - **Both**
8. In the **Packet Count** field, select the number of packets that you want to capture, from 1 to 8,000.

Packet capture stops when the selected number of packets has been captured.
9. Select **Add**.

The capture configuration for the selected port is displayed on the right.
10. Repeat [step 5](#) through [step 9](#) as needed to configure PCAPs for more ports.

PCAP configuration is supported for a maximum of 10 ports for the selected device.
11. Select **Save**.

The Packet Capture page displays running PCAPs and PCAP results.

Stop a PCAP

For SLX devices, you must stop the current uncompleted PCAP before you can begin a new one. 9920 devices support up to 10 running PCAPs before you need to stop one, although you do not need to reach the limit of 10 before stopping a PCAP.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. Select **Capture Packets** from the Device Actions menu.
4. Select **Configure Port Capture** or **+ (Enable PCAP)** as appropriate.
5. From the Packet Capture actions, select **Stop**.
6. Select **Delete** (🗑) for the packet capture you wish to stop.
The packet is removed from the list.

Download a PCAP File

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. Select **Capture Packets** from the Device Actions menu.
4. In the **Packet Capture** page, select **Download PCAP File** (⬇) for the PCAP file you want to download.
5. In the **File Name** field, provide a name for the file.
This step allows you to provide a user-friendly file name.
6. In the **Host IP** field, enter the IP address of a device that is accessible from the SLX device.
7. In the **User Name** and **Password** fields, provide the device credentials.
8. In the **Path** field, provide the download file path.
9. Select **Save**.
The PCAP file is downloaded to the specified destination.

Delete a PCAP File

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. Select **Capture Packets** from the Device Actions menu.
4. To delete multiple PCAP files, do the following:
 - a. Select the check boxes for the PCAP files you want to delete.
 - b. Select **Delete**.
The selected PCAP files are deleted.

- To delete a PCAP file, select **Delete PCAP File** () for the PCAP file you want to delete.

The selected PCAP file is deleted.

Clear Counters

You can clear counters for Extreme 9920, MLX, and SLX devices.

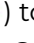
About This Task

Counters track the number packets. Counters increase over time and you can clear them as needed. For some devices, XCO supports specific object level counter clear and for some devices supports all object clear counter.

Table 12: Support for clearing counters

	9920	SLX	MLX
Clear all counters	<ul style="list-style-type: none"> • Interface • Match • Egress group • Ingress group • Egress • Ingress policy • Egress policy • Transport tunnel • Tunnel encapsulation 	<ul style="list-style-type: none"> • Interface • Match 	<ul style="list-style-type: none"> • Interface • Match
Clear specific counters	<ul style="list-style-type: none"> • Interface • Match • Egress group • Ingress group • Egress • Ingress policy • Egress policy • Transport tunnel • Tunnel encapsulation 	<ul style="list-style-type: none"> • Interface • Match 	<ul style="list-style-type: none"> • Interface

Procedure

- In the Navigation menu, select **Device Inventory**.
- In the **Devices** page, click anywhere in the required device row except the Actions column () to proceed to the device Overview page.
- Select **Clear Counters** from the Device Actions menu.

- To clear all counters of one object, select the check box for that type.

Selected

All Interfaces x

▼ Matches

▼ Interfaces

- To clear specific counters of one type, expand the type and select the check boxes for the counters.

Selected

3 Interfaces x

▼ Matches

▲ Interfaces

🔍 Search Interfaces

Ethernet 0/1

Ethernet 0/2

Ethernet 0/3

Ethernet 0/4

- Select **Clear**.
Reports in the dashboards are updated to reflect your selections.

View Logs

Procedure

- In the Navigation menu, select **Device Inventory**.
- In the **Devices** page, click anywhere in the required device row except the Actions column (☰) to proceed to the device Overview page.
- In the Device Actions menu, select **View Logs** to view the device specific logs.

Delete a Device from the Device Overview Page

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Actions menu, select **Delete**.

Policies and Configuration (Packet Broker Mode)

The **Policy Configuration** page in the user interface allows you to view, configure, or update existing device configurations. You can select the existing ingress groups, ingress policies, and egress-groups or create new objects and link them to create a complete service chain.

Policies

A policy represents the route-map or listener policy on the device. A policy consists of matches and actions to be applied on packets.

XCO supports the following policy types:

- Ingress Policy (route-map)
- Egress Policy (9920 only)

Ingress Policy

An ingress policy specifies the actions to be taken at ingress on packets entering the devices.

Egress Policy

An egress policy optionally specifies the actions to be applied on selected packets exiting the 9920 devices. You must configure a match ACL at a minimum.

Create an Egress Policy for a Device

An egress policy (or listener policy) defines the actions to apply to outbound packets.

Before You Begin

- Create a policy rule match to associate with the policy. For more information, see [Change a Policy Rule Match for a Device](#) on page 52.
- An ACL bound to an egress policy can be modified.
- An egress policy bound to an egress can be modified.

About This Task

Take the following steps to define the criteria for a policy. Each set of criteria is a rule. A policy can contain multiple rules.



Note

This topic applies only to Extreme 9920 devices.

Listener policy byte count is incorrect when truncation is enabled. The byte count for truncated packets is the actual byte count seen by the egress ACL before truncation.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policies > Add Policy**.
4. In the **Name** field, enter a unique name for the policy.
An egress policy cannot have the same name as another egress policy or the reserved keyword `all`.
5. In the **Policy Type** field, select **Egress Policy**.
6. Select the **Sequence** in which to apply the rule.
7. In the **Matches** field, select a policy rule match.
 - If you did not create a policy rule match, select **Create Match** to create the match now.
 - For a policy, you can select three rule matches of different types: 1 v4, 1 v6, and 1 I2.
 - You cannot use the same policy rule match multiple times in a policy.
8. In the **Packet Slicing** field, enter a value to represent the maximum packet size after slicing.
The final packet size will be less than or equal to this value.
9. In the **Header Stripping** field, select one or more tags to strip: 802.1BR, VLAN, or VN (Virtual NIC).
The 802.1BR and VN tags cannot coexist in the same policy rule action.
10. In the **VLAN** field, select the VLAN ID to change the VLAN tag in the egress packet.
11. To remove the outermost tunnel headers from the packet, select the **Decap** check box.
12. To prevent the rule from being used in the policy, select the **Deny** check box.



Tip

This option prevents the rule from being used, but does not delete the configuration of the rule. The rule is skipped and is not used to drop a packet. You can reinstate the rule later without having to reconfigure it.

13. Select **Add Rule**.

The rule parameters appear in the pane on the right.

14. Repeat step 7 through step 13 until you have added all the rules you need.
15. Select **Create**.

Create an Ingress Policy for a Device

An ingress policy (or route map) defines the actions to apply to inbound packets.

Before You Begin

Create a policy rule match to associate with the policy. For more information, see [Change a Policy Rule Match for a Device](#) on page 52.

Create an egress group to associate with the policy. For more information, see [Create an Egress Group](#) on page 58.

About This Task

Take the following steps to define the criteria for a policy. Each set of criteria is a rule. A policy can contain multiple rules.


Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policies > Add Policy**.
4. In the **Name** field, enter a unique name for the policy.
An ingress policy cannot have the same name as another ingress policy or the reserved keyword `all`.
5. In the **Policy Type** field, select **Ingress Policy**.
6. Select the **Sequence** in which to apply the rule.
7. In the **Matches** field, select a policy rule.



Note

- For a policy, you can select three rule matches of different types: 1 v4, 1 v6, and 1 L2.
- If you did not create a policy rule match, select **Create Match** to create the match now.
- You cannot use the same policy rule match multiple times in a policy.
- For SLX devices, you can select only one rule match type (v4, v6, L2, or UDA) per rule.
- For MLX devices, you cannot select L2 and UDA match in the same rule.
- (MLX only) Do not apply an L2 rule match and a UDA rule match in the same policy. Doing so causes the related ingress group to fail.
- (MLX only) If you add a UDA rule match to a policy that is associated with an ingress group, ensure that you first apply the associated UDA profile to that group. For more information, see [Create an Ingress Group for an SLX or MLX Device](#) on page 55.

8. (9920 only) In the **QoS** field, select the required QoS.
For more information, see [Quality of Service](#) on page 66.
 9. In the **Egress Group** field, select the group to associate with the policy.
If you did not create an Egress Group, select **Create Egress Group**. For more information, see [Create an Egress Group](#) on page 58.
 10. (MLX and 9920 only) In the **Packet Slicing** field, enter a value to represent the maximum packet size after slicing.
The final packet size will be less than or equal to this value.
 11. (SLX only) In the **Truncation Profile** field, select a profile that you created for a port or a port channel.
For more information, see [Create a Truncation Profile for an SLX Device](#) on page 71.
 12. (9920 only) In the Advance Scope section, select one of the following:
 - Decap** to remove the outermost tunnel headers from the packet
 - Scope Shift** to move the ACL scope for matching from the outer headers to the inner headers of a tunneled packet
 - None** to perform neither action
 13. (9920 only) To prevent the rule from being used in the policy, select the **Deny** check box.
-  **Tip**
This option prevents the rule from being used, but does not delete the configuration of the rule. The rule is skipped and is not used to drop a packet. You can reinstate the rule later without having to reconfigure it.
14. Select **Add Rule**.
The rule parameters appear in the pane on the right.
 15. Repeat step 7 through step 14 until you have added all the rules you need.
 16. Select **Create**.

Change a Policy for a Device

You can add, change, or delete one or more rules or actions in a policy.

About This Task

You can change a policy for a specific device or change a policy in the library. To change a policy in the library, see [Change a Policy in the Library](#) on page 104.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policies**.
4. Select the policy that you want to change.

5. Follow the instructions in [Create an Ingress Policy for a Device](#) on page 46 or [Create an Egress Policy for a Device](#) on page 44 to add, change, or remove rules or actions in the policy.

**Tip**

(9920 only) To reinstate a rule that is not included in the policy (the **Deny** field is selected), clear the **Deny** field.

*View the Policy Configuration***About This Task****Procedure**

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select **Policies and Configuration**.

The Policy Configuration page displays the existing Ingress Group, Ingress Policy, and Egress Group information for the device.

*Delete a Policy from a Device***About This Task**

You can delete a policy from a device or from the library. To delete a policy from the library, see [Delete a Policy in the Library](#) on page 105.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policies**.
4. Select **Delete Policy** for the policy that you want to delete.
5. Remove the policy from any associated ingress group. For more information, see [Change an Ingress Group](#) on page 57.

Policy Rule Matches

A policy rule match represents IPv4, IPv6, L2, or UDA configuration on the device.

*Create a Policy Rule Match for a Device***About This Task**

When you create a policy rule match, you select all parts of a packet header that you want to target and then select the action to perform on the targeted items. These

selections are the rules in your match. The match can then be associated with ingress or egress policies. A policy rule match can contain one or more rules.



Note

A policy rule match is a device-specific feature. If you have UDAs configured for a device, UDA-related fields are displayed in the Create Match page. These fields are not described in this procedure.

XCO supports a maximum of 6000 IPv4, 2000 IPv6, and 1500 L2/MAC matches for 9920.

To create a policy rule match in the library, see [Create a Policy Rule Match in the Library](#) on page 99.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policy Rule Matches > Add Policy Rule Match**.
4. In the **Name** field, enter a unique name for the match.
 - Alphanumeric characters, dashes, and underscores are allowed in the **Name** field.
 - The name, `all` is a reserved keyword on 9920 and cannot be used.
5. In the **Type** field, select whether the match applies to IPv4, IPv6, L2, or UDA. If you selected UDA on an SLX device, proceed to the next step. Else, go to step 7.
6. In the **UDA** field, select a profile.
7. (SLX only) In the **Sub Type** field, select the appropriate match.
 - **Standard**: Matches the source address information
 - **Extended**: Matches the source and destination address information
8. In the Match section, complete the applicable fields to identify the packets of interest.



Note

All fields are not mandatory. You can leave the fields blank unless noted.

The items in this section vary by your selection in the **Protocol** field. The following list describes all possible selections.

- **Protocol**: The protocol that you want to target. If the protocol you want is not in the list, select **None** and provide the ID of the protocol you want in the **Protocol ID** field. Every protocol has a numeric value that is defined by the IETF.
- **Sequence**: The order in which this rule is performed in the match.
- **Protocol ID**: The ID of a protocol that you want to target. Use only when the protocol you want is not available in the **Protocol** field.
- **Source IP**: The IPv4 or IPv6 address of the device that sends the packets.
- **Source Mask**: The mask for the source IP address, in the following format: 255.255.255.255 or ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff.

- **Destination IP:** The IPv4 or IPv6 address of the device that is to receive the packets.
- **Destination Mask:** The mask for the destination IP address, in the following format: 255.255.255.255 or ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff.
- **Source Mac:** The MAC address of the device that sends the packets, in the following format: 1111.1111.1111 or 11:11:11:11:11:11. Any alpha characters in the address must be lowercase.
- **Source Mac Mask:** The mask for the source MAC address, in the following format: ffff.ffff.ffff or ff:ff:ff:ff:ff:ff. Any alpha characters in the mask must be lowercase.
- **Destination Mac:** The MAC address of the device that is to receive the packets, in the following format: 1111.1111.1111 or 11:11:11:11:11:11. Any alpha characters in the address must be lowercase.
- **Destination Mac Mask:** The mask for the destination MAC address, in the following format: ffff.ffff.ffff or ff:ff:ff:ff:ff:ff. Any alpha characters in the mask must be lowercase.
- **Source Port:** The port through which packets enter the device.
- **Source Port End:** The last port in the range of ports through which packets enter the device.
- **Destination Port:** The port through which packets leave the device. Valid values range from 1 through 65535.
- **Destination Port End:** The last port in the range of ports through which packets leave the device. Valid values range from 1 through 65535.
- **IP Payload Length:** The length of the IP packets that you want to target, or the size of the IP payload. Valid values range from 64 through 9000.
- **IP Payload Length End:** The last acceptable value of the IP payload. Valid values range from 65 through 9000.
- **DSCP:** The value of the Differentiated Services Code Point in the Type of Service field in the header. Valid values range from 0 through 63.
- **VLAN:** The VLAN ID. The valid value ranges are as follows:
 - 9920: 0 through 4095
 - SLX and MLX: 0 through 4091
- **EtherType:** Identifies the protocol that is encapsulated in the payload. For example, the EtherType value for IPv4 is 0x0800. Valid values range from 1536 through 65536 (numerical), or 0x0600 through 0xffff (hexadecimal), or are one of the following: ARP, IPv4, or IPv6.
- **PCP:** The Priority Code Point, a 3-bit field in a VLAN header. Valid values range from 0 through 7.
- **Tunnel ID:** The ID number of the tunnel. Valid values range from 1 through 16777215.

- **MATCH0, MATCH1, MATCH2, MATCH3:** Specifies the UDA Hexadecimal. SLX presents these as specific header fields such as `NEXT_HEADER`.



Note

- MLX UDA requires a match and mask for all fields.
- Use a mask of all zeros to make the any value for a field.

- **MASK0, MASK1, MASK2, MASK3:** Specifies the UDA Hexadecimal value used to mask the MATCH values. Use 0 bits for any value. A bit value of 1 must be matched.

9. In the Fragmentation sub-section, select one of the following.

The items in this section vary by your selection in the **Type**, **Sub Type** and **Protocol** fields. The following list describes all possible selections.

- **Fragmented:** Targets target fragmented packets.
- **Non Fragmented:** Targets non-fragmented packets.
- **None:** Targets packets in which the DF (Don't Fragment) flag is set in the IP header.

10. In the Options sub-section, select one or more of the following:

The items in this section vary by your selection in the **Type**, **Sub Type** and **Protocol** fields, in particular selection of a Layer4 protocol such as UDP, TCP, or STCP. The following list describes all possible selections.

- **Acknowledgment:** Targets packets in which the ACK flag is set in the TCP header.
- **Congestion:** Targets packets in which the CWR flag is set in the TCP header.
- **ECN-Echo:** Targets packets in which the ECE flag is set in the TCP header.
- **Last Packet:** Targets packets in which the FIN flag is set in the TCP header.
- **Push:** Targets packets in which the PSH flag is set in the TCP header.
- **Reset:** Targets packets in which the RST flag is set in the TCP header.
- **Synchronize:** Targets packets in which the SYN flag is set in the TCP header.
- **Urgent:** Targets packets in which the URG flag is set in the TCP header.

11. In the Action section, select one or more actions to perform on the targeted items.

- **Drop** to drop the packet.
- **Count** to keep track of the number of packets that match the policy rule.
- **Log** to add the transaction to the XCO log.

12. Select **Add**.

The match parameters (the new rule) appear in the pane on the right.

13. Repeat steps 8 through 12 until you have added all the rules you need.

14. Select **Save**.


Change a Policy Rule Match for a Device

You can add, change, or delete one or more rules in a policy rule match.

About This Task

You can change a policy rule match for a specific device or change a match in the library. To change a match in the library, see [Change a Policy Rule Match in the Library](#) on page 102.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policy Rule Matches**.
4. Select **Edit** () from the Actions column (⋮) for the policy rule match that you want to change.
5. Follow the instructions in [Create a Policy Rule Match for a Device](#) on page 48 to add, change, or remove rules in the match.
6. Select **Update** to save the configuration changes.

Delete a Policy Rule Match from a Device

You can delete a policy rule match from a device.


Before You Begin

You cannot delete a policy rule that is attached to a device.

About This Task

To delete a policy rule match from the library, see [Delete a Policy Rule Match from the Library](#) on page 103.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Policy Rule Matches**.
4. Select **Delete** () from the Actions column (⋮) for the match you want to delete.

UDA Profiles

The UDA profile consists of offset values. The UDA profile must be attached to the ingress group to apply it to all interfaces.

- A maximum of four parameters per profile are supported.
- Only multiples of four are supported as offset values for MLX devices.
- For SLX devices, the profiles are updated in the UDA match.

Create an MLX UDA Profile for a Device

About This Task

To create an MLX UDA profile in the library, see [Create an MLX UDA Profile in the Library](#) on page 106.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > UDA Profiles > Add UDA Profile**
4. In the **Name** field, enter a unique name for the UDA profile.
5. In the four **Offset** fields, select the appropriate offset values.
An offset indicates the index of the received packet. For example, an offset of 0 indicates the first byte of the received packet.
6. Select **Save**.

Create an SLX UDA Profile for a Device

A UDA profile can be associated with a UDA match.

About This Task

To create an SLX UDA profile in the library, see [Create an SLX UDA Profile in the Library](#) on page 106.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > UDA Profiles > Add UDA Profile**
4. In the **Name** field, enter a unique name for the UDA profile.
5. Define the header fields that are required for a match.
The header fields you select constitute the header stack. As you select header types and header fields, additional header selections become available. The additional selections vary based on your header choices.
 - a. In the **Header 0 Ethernet - Ethernet** row, select the field that is required for a match and then click **+** to add your selection.
 - b. In the **Header 1** row, select the type and field that are required for a match and then click **+** to add your selection.
Your selections determine whether a Header 2 row is displayed.
 - c. Make selections in the Header 2 row and in all subsequent rows until no more rows are available or until your header stack is complete.
A maximum of 4 Headers are supported in a UDA profile.
6. Select **Save**.


Change a UDA Profile for a Device

You can change the parameters of a user-defined access list (UDA) profile.

About This Task

To change a UDA profile in the library, see [Change a UDA Profile in the Library](#) on page 107.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > UDA Profiles**.
4. In the UDA Profiles page, select **Edit** () from the Actions column (⋮) for the UDA that you want to change.
5. To change a UDA Profile for a device, take the following steps.
6. Follow the instructions for the type of UDA you are changing.
 - [Create an SLX UDA Profile for a Device](#) on page 53
 - [Create an MLX UDA Profile for a Device](#) on page 53
7. Select **Save**.

Delete a UDA Profile from a Device

You can delete a user-defined access list (UDA) profile from the library or device inventory page.


About This Task

To delete a UDA profile in the library, see [Delete a UDA Profile in the Library](#) on page 107.

Before You Begin

You cannot delete a UDA profile that is attached to any ingress-group.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > UDA Profiles**.
4. In the UDA Profile page, select **Delete** () from the Actions column (⋮) for the UDA Profile you want to delete.

Ingress Groups

Ingress policies are used to perform actions on packets at ingress. Ingress groups classify the packets received on devices.

Create an Ingress Group for an SLX or MLX Device

An ingress group is a set of ports and port channels on which monitored traffic is received.

Before You Begin

If necessary, create the port channel to associate with the ingress group. For more information, see [Create a Port Channel](#) on page 68.

If necessary, create the ingress policy to associate with the ingress group. For more information, see [Create an Ingress Policy for a Device](#) on page 46.

If necessary, create a UDA profile to associate with the ingress group. For more information, see [Create an MLX UDA Profile in the Library](#) on page 106.

About This Task

Ingress groups classify and apply policies on monitored traffic. After you create an ingress group, the group can be associated with an ingress policy.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Ingress Groups > Add Ingress Group**.
4. In the **Name** field, enter a name for the group.
5. In the **Ports/Port Channels** field, select at least one port or port channel for the group.
6. In the **Policy Name** field, select the ingress policy to associate with the group.
7. In the **UDA Profile** field, select a UDA profile to associate with the group.
You must select a profile if the policy (in the **Policy Name** field) contains a UDA match. If you do not select a profile, your ingress group configuration will fail.
8. Select **Create**.

Create an Ingress Group for a 9920 Device

An ingress group is a set of ports, port channels, and tunnels on which monitored traffic is received.

Before You Begin

If necessary, create the port channel to associate with the ingress group. For more information, see [Create a Port Channel](#) on page 68.

If necessary, create the ingress policy to associate with the ingress group. For more information, see [Create an Ingress Policy for a Device](#) on page 46.

If necessary, create a mirror for the outer tunnel. For more information, see [Configure a Traffic Mirror for 9920 Devices](#) on page 62.

About This Task

Ingress groups classify and apply policies on monitored traffic. After you create an ingress group, the group can be associated with an ingress policy.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Ingress Groups > Add Ingress Group**.
4. In the **Name** field, enter a name for the group.
5. In the **Ports/Port Channels** field, select at least one port or port channel for the group.
6. (Optional) For single tunnel encapsulation, do the following to configure the Inner Tunnel:
 - a. Expand the **Inner Tunnel** section.
 - b. In the **Tunnel Type** field, select the type of tunnel for the incoming traffic.
 - **GRE**
 - **GTPU**
 - **VXLAN**
 - **NVGRE**
 - **IPIP**
 - c. (Optional) In the **Tunnel ID** field, select or enter a value that represents the tunnel ID.

This field is not applicable for GRE and IPIP tunnels.
 - d. (Optional) Complete the applicable processing and filter options for the selected protocol.
 - **Destination IP**: Specifies the destination IP address
 - **Destination Prefix**: Specifies the destination prefix
 - **Source IP**: Specifies the source IP address
 - **Source Prefix**: Specifies the source prefix
 - e. In the Advance Scope section, select one of the following actions to apply to the incoming traffic.
 - **Decap** to remove the outer tunnel headers from the packet
 - **Scope Shift** to move the ACL scope for matching from the outer headers to the inner headers of a tunneled packet
 - **None** to perform neither action

7. (Optional) For packets with two sets of tunnel headers before the innermost packet, for example, a VXLAN tunnel wrapped around a GTPu tunneled packet, do the following to configure the Outer Tunnel.
 - a. Expand the **Outer Tunnel** section.
 - b. In the **Outer Tunnel Type** field, select the type of tunnel for the incoming traffic.
 - **None**
 - **VXLAN**
 - **MPLS**

A maximum of five MPLS header removal is supported. The packets with more than five MPLS headers are dropped.
 - c. Complete the applicable filter options for the outer tunnel headers for the selected protocol.
 - **Label:** Filters on the last MPLS label present in a five label stack.
 - **Traffic Class:** Filters on the Traffic Class field of the last MPLS label present in a five label stack.
 - **Time To Live (TTL):** Filters on the Time To Live field in the last MPLS label present in a five label stack.
 - **Outer Tunnel ID:** Filters on the VXLAN tunnel ID field.
 - **Outer Destination IP:** Specifies the destination IPv4 address or network..
 - **Outer Destination Prefix:** Specifies the destination prefix if filtering on a range of hosts.
 - **Outer Source IP:** Specifies the source IPv4 address or network.
 - **Outer Source Prefix:** Specifies the source prefix if filtering on a range of hosts.

The packets that do not match the selected filter options are dropped.
 - d. (Optional) In the **Mirror** field, select the mirror action to forward a copy of the entire packet to the configured mirrored port.

For more information, see [Configure a Traffic Mirror for 9920 Devices](#) on page 62.
8. In the **Policy Name** field, select the ingress policy to associate with the group.
9. Select **Create**.

Change an Ingress Group

You can add, change, or delete the parameters of an ingress group.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (☰) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Ingress Groups**.
4. In the list of ingress groups, select the group that you want to change.
5. In the Actions column, select **Edit** (✎) for the ingress group.

6. Follow the instructions in [Create an Ingress Group for an SLX or MLX Device](#) on page 55 and [Create an Ingress Group for a 9920 Device](#) on page 55 to add, change, or delete the parameters in the group.

Delete an Ingress Group

You can delete an ingress group from a device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Ingress Groups**.
4. In the Actions column, select **Delete** for the group that you want to delete.

Egress-Group

An egress-group represents one or more interfaces for replicating the traffic for the given policy.

An egress-group connects an ingress-policy and the egress to define how traffic is forwarded to end devices.

Create an Egress Group

An egress group is a set of egresses (9920) or a set of interfaces or port-channels (SLX/MLX).

Before You Begin

Create the egress (9920) or port-channels (SLX/MLX) to associate with the egress group. For more information, see [Create an Egress for Devices](#) on page 60 and [Create a Port Channel](#) on page 68.

About This Task

When you create an egress group, you assign a name and associate at least one egress (9920) or port/port channel (SLX/MLX). An egress associates an egress port (or port channel) with an egress policy for 9920.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Egress Groups > Add Egress Group**.

4. To create an egress group for SLX and MLX devices, take the following steps.
 - a. In the **nHop Type** field, select the next hop domain type: **TVF** (transparent VLAN flooding) or **VLAN** (MLX only).
 - b. In the **nHop Value** field, provide the VLAN ID (MLX only) or TVF ID.
The valid value ranges for VLAN and TVF are as follows:
 - MLX: VLAN is 1-4090 and TVF is 1-2016
 - SLX: TVF is 1-4095
 - c. Select the required **Ports/PortChannels**.
 - d. Select **Create**.
5. To create an egress group for 9920 devices, take the following steps.
 - a. In the **Name** field, enter a name for the group.
An egress group cannot have the same name as an egress.
 - b. In the **Egress List** field, select at least one egress to associate with the group.
 - c. (Optional) Select **Create Egress** to create an egress to associate with the egress group, if required.
For more information, see [Create an Egress for Devices](#) on page 60.
 - d. Select **Create**.

Change an Egress Group

You can add or delete egress in an egress group.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Egress Groups**.
4. In the Actions column, select **Edit** for the egress group that you want to change.
5. In the **Egress List** (9920) or **Interface/Port Channel** (SLX/MLX) field, select (or delete) at least one item.
For more information, see [Create an Egress Group](#) on page 58.
6. Select **Save**.

Delete an Egress Group

You can delete an egress group from a device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Egress Groups**.
4. In the Actions column, select **Delete** for the group that you want to delete.

Egress

An egress defines an interface or a list of interfaces to be used for tool ports.

You can create an egress and combine it with various policies to perform additional processing actions to packets leaving the egress interfaces.

For MLX and SLX devices, the egresses are created internally when the egress group is created, and the egress group lists the ports or port-channels along with TVF or VLAN.

Create an Egress for Devices

Egress is a combination of ports, port channels, precedence, and an associated egress policy.

Before You Begin

If necessary, create a port channel. For more information, see [Change a Port Channel](#) on page 69.

(9920 only) If necessary, create an egress policy. For more information, see [Create an Egress Policy for a Device](#) on page 44.

About This Task



Note

For MLX and SLX devices, the egresses are created internally when the egress group is created, and the egress group lists the ports or port-channels along with TVF or VLAN.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Egresses > Add Egress**.
4. In the **Name** field, enter a name.
The egress cannot have the same name as an egress group.
5. In the **Port/Port Channel** field, select an egress port or port channel.
6. In the **Precedence** field, select the order of precedence for the port or port channel.
The precedence indicates the priority given to the port or port channel. The lower the number, the higher the priority.
7. Select **Add Port/Precedence (+)** to add your selections.
8. (9920 only) In the **Egress Policy** field, select the policy to associate with the egress.
9. Select **Save**.

Change an Egress Configuration

You can change the parameters of the egress configuration for a 9920 device.

Before You Begin

The egress configuration is view-only for SLX and MLX devices.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select **Policies and Configuration > Egresses**.
4. In the **Actions** column, select **Edit** for the item that you want to change.
5. Complete the fields as described in [Create an Egress for Devices](#) on page 60.

Delete an Egress Configuration

You can delete the egress configuration from a 9920 device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Egresses**.
4. In the Actions column, select **Delete** for the item that you want to delete.

Mirrors

XCO supports packet mirroring only for the 9920 devices.

Packet mirroring mirrors the whole frame to another egress port. For a frame without the outer tag, the header is terminated and the frame is subjected to further processing, based on ingress and egress configuration.

When mirroring is enabled, one copy of the whole frame is subjected to normal processing where the header is terminated and subjected to regular ingress or egress processing.

Another copy of the frame is mirrored with egress port without any header termination.

The filters for frame that are configured using ingress-group can be applied per port. If the frame does not match the filter, it is dropped.



Note

- Only one mirror destination port is supported.
- You can use the ingress-group to enable mirroring for outer MPLS-SR and outer VXLAN termination. MPLS-SR packets that match the filters are sent to the egress port based on the configured mirror.

Configure a Traffic Mirror for 9920 Devices

You can mirror traffic to a mirror port interface.

About This Task

The mirror is used in the outer tunnel configuration for an ingress group. This process ensures that the designated mirroring destination receives the same traffic as the egress port.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Mirrors > Add Mirror**.
4. In the **Name** field, enter a name for the mirror.
5. In the **Description** field, enter the description for the mirror.
6. In the **Port** field, select the mirroring destination port.
7. Select **Save**.

Change a Mirror Configuration

You can change the parameters of the configuration.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select **Policies and Configuration > Mirror**.
4. Select the item that you want to change.
5. Complete the fields as described in [Configure a Traffic Mirror for 9920 Devices](#) on page 62.

Delete a Mirror Configuration

You can delete the configuration from a device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Mirror**.
4. In the Actions column, select **Delete** for the item that you want to delete.

Tunnels

XCO supports GRE tunnel encapsulation and termination on 9920 devices for forwarding selected traffic from a local 9920 device to a remote 9920 device through a routed network. These packets are sent to the specified interface on 9920 compared

to other tunneled traffic that is mirrored or copied, but not addressed explicitly to the 9920.

**Note**

- XCO supports tunnel configuration only for the 9920 devices.
- As a best practice, configure static ARP entries on the routers for the connected 9920 device.

Encapsulation

Generic Routing Encapsulation (GRE) headers provide a private secure path for transporting packets.

The following information is required for tunnel creation or encapsulation:

- Source MAC address
- Source IPv4 address
- Destination MAC address
- VLAN ID
- Destination IPv4 address

The destination IP address must be in the network of the remote router.

Tunnel Termination

XCO decapsulates packets based on the configured parameters. The following information is required for tunnel termination:

- Source IPv4 address
- Source Prefix
- Destination IP address
- Destination Prefix

Tunnel Termination Flow

Tunnel termination can be configured for received L2 or L3 packets.

You can configure settings to apply tunnel termination to received packets, either tunneled (both L2 and L3) or non-tunneled. Tunnel termination is performed at either ingress or egress depending on the policy configuration.

L2 tunnel termination flow is as follows:

1. The outer tunnel of L2 tunneled packets is removed.
2. The current position is shifted to the start of the inner L2 header.

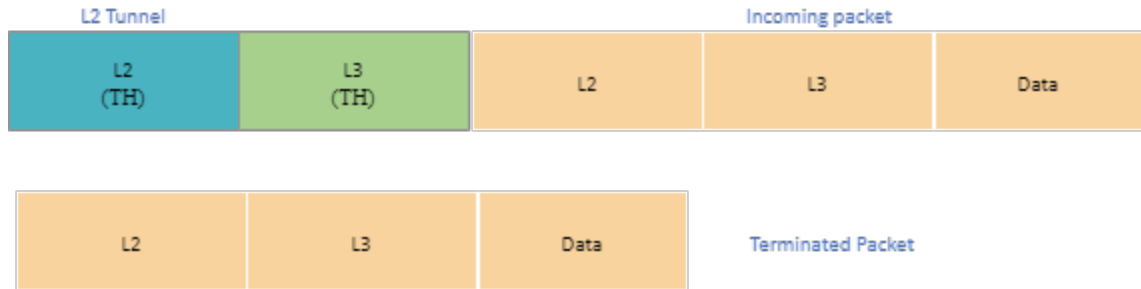


Figure 4: L2 tunnel termination

L3 tunnel termination flow is as follows:

1. The L2 header is retrieved from the L2 outer header because L3 tunneled-packet inner headers do not have the L2 header.
2. The L3 outer header is stripped.




Figure 5: L3 tunnel termination

Create a Tunnel

You can configure transport tunnel termination and encapsulation for a device.

About This Task

You can associate transport tunnel termination with an ingress group and then associate that group with an ingress policy.



Note

This feature applies to Extreme 9920 devices only.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Tunnels > Add Tunnel**.
4. In the **Name** field, enter a name for the tunnel.
5. In the **Type** field, select **Termination** or **Encapsulation**.

6. In the **Tunnel Type** field, select one of the following.

The options vary by the type you selected in step 5.

GRE (Generic Routing Encapsulation). This type of tunnel encapsulates (or wraps) packets that use a certain protocol inside packets that use a different protocol.

ERSPAN (Encapsulated Remote Switched Port Analyzer): This type of tunnel mirrors traffic from source ports for delivery to destination ports on a different device.

7. In the **Source IP** field, enter the IPv4 or IPv6 address of the device that sends the packets.
8. In the **Destination IP** field, enter the IPv4 or IPv6 address of the device that is to receive the packets.
9. Complete the following fields.

The fields vary by the type you selected in step 5.

Source MAC. The MAC address of the device that sends the packets.

Destination MAC. The MAC address of the gateway router.

VLAN Tag. A numeric string that identifies which VLAN a packet belongs to.

VLAN PCP. The Priority Code Point, a 3-bit field in the VLAN header.

Egress. The egress to associate with the tunnel.

Source Prefix. The prefix of the IP address of the device that sends the packets, in CIDR notation format.

Destination Prefix. The prefix of the IP address of the device that receives the packets, in CIDR notation format.

Ingress Groups. The ingress group to associate with the tunnel.

10. Select **Save**.

Change a Tunnel

You can change the tunnel configuration for a device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Tunnels**.
4. Select the tunnel that you want to change.
5. Follow the steps in [Create a Tunnel](#) on page 64 to change the tunnel configuration.

Delete a Tunnel

You can delete tunnel configuration from a device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Tunnels**.

4. Select **Delete** for the tunnel that you want to delete.

Quality of Service

XCO supports Quality of Service (QoS) configuration on 9920 devices.

QoS provides preferential treatment to specific traffic that is received on multiple ingress interfaces or Test Access Points (TAPs). QoS enables a networking device prioritize critical traffic streams and provides dedicated bandwidth for effective delivery.

QoS aggregates, filters, and forwards traffic to a monitoring tool on an egress interface or egress group. The forwarding decision is based on the access control lists (ACLs) and route maps applied on the aggregated logical interface or port channel.

QoS can selectively drop the low priority traffic streams to allow high priority traffic to pass through. QoS manages traffic delivery using queues, buffers, and schedulers for maximum throughput.

QoS supports eight queues per egress port on a device. The highest queue priority is q7 and q0 is the lowest queue priority.

The configured QoS can be used in policy configuration and rule matches. For more information, see [Create an Ingress Policy for a Device](#) on page 46.


Add a QoS

Procedure

1. In the navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > QoS > Add QoS**.
4. In the **Name** field, enter a name.
5. In the **Description** field, enter the description.
6. In the **Queue** field, select the queue priority.
The highest queue priority is q7 and q0 is the lowest queue priority.
7. Select **Save**.

Change a QoS

Procedure

1. In the navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (⋮) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > QoS**.
4. In the QoS page, select **Edit** () from the Actions column (⋮) for the QoS you want to change.
5. Follow the instructions in [Add a QoS](#) on page 66 to change the QoS parameters.

6. Select **Save**.

Delete a QoS

Procedure

1. In the navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > QoS**.
4. In the QoS page, select **Delete** from the Actions column (•••) for the QoS you want to delete.

Port Channels

Port channels, also called Link Aggregation Groups (LAG), are used for load balancing traffic among ports.

Link Aggregation

Link Aggregation (LAG) bundles multiple physical Ethernet links into a single high-bandwidth port-channel for enhanced performance and redundancy.

LAG provides load balancing across physical interfaces and improves reliability. The port-channel stays operational as long as at least one physical interface within the port-channel is operational.

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE standards-based protocol that allows two partner systems to dynamically negotiate attributes of physical links between them to form logical links.

LACP enables devices to send Link Aggregation Control Protocol Data Units (LACPDU) to each other and establish link aggregation connections.

Static LAG

In static link aggregation, you can create a port-channel interface or LAG and add member interfaces manually.

In static link aggregation, Link Aggregation Control Protocol (LACP) packets are not exchanged between the partner systems. Aggregation and load-balancing of frames on static links is determined by the operational status and administrative state of the link.

Minimum Links

Minimum links define the operational state of a LAG interface. If the number of operationally up Ethernet ports are less than configured minimum links value, the LAG interface is considered operationally down. By default, minimum links value is set to 1. At least one member port must be up, for a LAG interface to be operationally up.

Create a Port Channel

Before You Begin

The MTU and egress configuration must be same on all ports prior to configuring a LAG. If the Egress configuration is different, take the following steps:

1. Remove the desired ports from the Egress they are associated with.
2. Create the LAG and add the LAG back into the appropriate Egress.

About This Task

After you create a port channel, it is available for selection when you create ingress group and egress.



Note

- The fields that are available for creating a port channel vary by the device type you are configuring.
- LACP LAG is supported for 9920.
- For SLX devices, static LAG type is selected by default.
- All configurations are reverted when a port channel deployment fails. However, a LAG is created and deleted immediately, and the events are captured in the device logs.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Port Channels**.
4. Select **Add Port Channel**.
5. (MLX only) In the **Port Channel Name** field, enter a unique name for the port channel.
6. (9920 and SLX) In the **Port Channel ID** field, enter a unique numeric ID.
No two port channels can have the same ID.
7. In the **Lag Speed** field, select the speed for the ports that you will select in step 9.
8. In the **Description** field, provide enough information to help you identify the port channel.
9. In the **Ports** field, select at least one port from the list.
The ports in the list will be of the speed that you selected in step 7. A port can be a member of only one port channel. Ports that are not in the list are either already added to another port channel or are operating at a speed that is different from the selection in step 7.
10. (SLX and 9920 only) In the **MTU** field, enter the maximum transmission unit for packets that pass through the ports in the channel.
(SLX only) If the Port Speed is configured as auto, but not connected, the Port Speed must be updated manually to refresh the list of ports.
11. (MLX only) In the **Primary Port** field, select one of the member ports.
12. In the **Minimum Link** field, select or enter the minimum number of interfaces that the port channel requires to be active.

13. (9920 only) In the **Load Balance Algorithm** field, select a load-balancing method or select **None**.
 - src-dst-ip-l4-port**: The source and destination IP Layer 4 ports method is the default load-balancing method.
 - src-dst-ip-l4-port-tid**: The source and destination IP Layer 4 ports method with tunnel ID.
14. Select **Enable** to change the port channel admin status to Up.

When you select this field, you initiate the **no shutdown** command on the device, which changes the admin and operating status. When the field is not selected, the **shutdown** command runs on the port channel and the admin status changes to Down.
15. Select the **Lag Type**.
 - **STATIC**
 - **LACP**
16. (MLX and SLX only) Select **Loopback** to configure the port channel as a loopback interface.

A loopback is a virtual interface that a device uses to communicate with itself. A loopback interface cannot be used as an egress interface.
17. Select **Save**.

Change a Port Channel

You can change the parameters of a port channel.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Port Channels**.
4. Select the object you want to change.
5. Follow the steps in [Create a Port Channel](#) on page 68 to change the channel parameters.

Few of the parameters are read-only and cannot be changed.

Delete a Port Channel

You can delete a port channel from a device.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Port Channels**.
4. Select **Delete** for the object that you want to delete.

Ports

XCO supports port or port-channel configurations.

Configure Port Properties

You can change several port properties, including description, port speed, MTU, Forward Error Correction (FEC), port breakout, header stripping, Link Fault Signaling (LFS), admin state.

About This Task



Note

(SLX only) If the port you are configuring is part of a port channel, do not change the **MTU** or the **Port Speed** values from the ExtremeCloud Orchestrator interface.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Ports**.
4. In the Actions column, select **Edit** for the port you want to configure.
5. In the **Description** field, provide new information.
6. In the **Port Speed** field, select a different speed.
7. In the **MTU** field, enter the maximum transmission unit for packets that pass through the port.
8. (MLX and SLX only) In the **Headers** field, select one or more types of header to strip.
9. (9920 only) In the **Breakout** field, select one of the following:
 - **4x10G**: Configures four 10G breakout interfaces on the port.
 - **4x25G**: Configures four 25G breakout interfaces on the port.
 - **None**



Note

Ports with auto speed configured and not connected, cannot have speed reconciled. Without the speed reconciliation, ports are not listed in the port-channel and are not allowed to breakout or un-breakout.

10. (SLX and 9920 only) In the **Forward Error Correction** field, select one of the following:
 - **Auto-negotiation**: Selects the appropriate algorithm automatically.
 - **FC-FEC**: An algorithm that corrects errors in a block of data, with lower latency than RS-FEC.
 - **RS-FEC**: An algorithm that corrects errors in a block of data, with better error correction than FC-FEC.
 - **Disabled**: Disables the FEC feature.

FEC corrects errors in data without the need for retransmission of the data. Port speed determines which FEC configuration is supported.

- For 100G ports, RS-FEC, Auto-negotiation, and Disabled are supported.
- For 25G ports, RS-FEC, FC-FEC, Auto-negotiation, and Disabled are supported.
- For 40G and 10G ports, only Disabled is supported.



Note

FEC can be updated only when the port is in shutdown state.

11. (9920 only) To enable communication between two Ethernet devices, select **Link Fault Signaling**.

Link Fault Signaling is a physical layer protocol that enables a port to report fault conditions on inbound and outbound ports.

12. Select **Enable** to change the port admin status to Up.

When you select this field, you initiate the **no shutdown** command on the device, which changes the admin status to Up. When the field is not selected, the **shutdown** command runs on the port and the admin status changes to Down.

13. (MLX only) Select the required **Port Type**.

- **INGRESS**
- **EGRESS**
- **SERVICE**

14. (MLX and SLX only) Select **Loopback** to configure the port as a loopback interface.

A loopback is a virtual interface that a device uses to communicate with itself. A loopback interface cannot be used as an egress interface.

15. Select **Save**.

Truncation Profile

A truncation profile is used for packet slicing in SLX devices. A maximum of four truncation profiles are supported for a device.

Create a Truncation Profile for an SLX Device

About This Task



Note

When a port is configured for truncation, it becomes a loopback port. When the truncation profile is deleted, the loopback mode is removed.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Truncation Profile > Add Truncation Profile**.

4. In the **Profile Name** field, enter a name for the truncation profile.
5. In the **Frame Size** field, enter a value to represent the maximum packet size after truncation.
The frame size must be a multiple of 16 and valid range is 64 to 9216.
6. In the **Ethernet Interface** field, select the interface or in the **Port Channels** field, select the port-channel as appropriate.
7. Select **Save**.

Change a Truncation Profile

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Truncation Profile**.
4. Complete the fields as described in [Create a Truncation Profile for an SLX Device](#) on page 71.
5. Save your changes.

Delete a Truncation Profile

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Device Config menu, select **Policies and Configuration > Truncation Profile**.
4. In the **Actions** column, select **Delete** for the truncation profile you want to delete.

TACACS+ Authentication

Terminal Access Controller Access-Control System Plus (TACACS+) is an external authentication server used for verifying user credentials. For more information, see [TACACS+ Settings](#) on page 92.

Configure TACACS+ Authentication for Device Access

For support of the TACACS+ servers you have configured, you must enable TACACS+ authentication.

About This Task

The default authentication value for Extreme 9920, SLX, and MLX devices is always local, so you explicitly change the authentication to TACACS+ when you add a TACACS+ server.

Procedure

1. In the Navigation menu, select **Device Inventory**.

2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. In the Devices Config menu, select **Policies and Configuration > Authentication**.
4. To change authentication from local to TACACS+ on the device, take the following steps.
 - a. In the Actions column, select **Edit**.
 - b. In the **Authentication Type** field, select **TACACS+**.
 - c. Select **Save**.
5. To change authentication from TACACS+ to local, take the following steps.
 - a. In the Actions column, select **Edit**.
 - b. In the **Authentication Type** field, select **Local**.
 - c. Select **Save**.

Slots

You can view and update the configuration of the slots for a selected MLX device.

Change MLX Slot Configuration

About This Task

All available slots are displayed in the XCO user interface in the **Policies and Configuration** page of the device detail view. For a selected packet processor of the selected slot, you can change the configuration for header stripping, packet slicing, and packet length match.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.

3. In the Device Config menu, select **Policies and Configuration > Slots**.

← MLXe_247_QA Device Actions

Overview

Policies and Configuration

- Policies
- Policy Rule Matches
- UDA Profiles
- Ingress Groups
- Egress Groups
- Egresses
- Port Channels
- Ports
- TACACS+
- Slots**

Slots

Showing 1 - 8 of 8 results Group By Page Size Latest as on 5:38:30 PM

Name	Actions
Slot 1	
Slot 2	
Slot 3	
Slot 4	
Slot 5	
Slot 6	
Slot 7	

< 1 >

4. Select **Edit** () from the Actions column () for the slot you want to update.

Edit Slot : Slot 1

Processor

Processor**

Header Stripping

Headers

Packet Slicing

Truncate Egress Ports Truncate Size

Packet Length Match

Minimum Length Maximum Length

Search Config

processor1

- Headers: BR802,NVGRE,VXLAN,VNTAG,BYPASS
- Truncate Egress Ports: ethernet 1/1,ethernet 1/2,ethernet 1/5,ethernet ...
- Truncate Size: 9000 Minimum Length: 67
- Maximum Length: 200

processor2

- Headers: NVGRE,VNTAG,VXLAN,BR802,BYPASS
- Truncate Egress Ports: ethernet 1/11 Truncate Size: 9197
- Minimum Length: 0 Maximum Length: 0

The right side of the page displays the current configuration for each processor in the device. You can add, edit, or delete one or more actions for each processor.

Parts Details (Packet Broker Mode)

You can view and download the device inventory.

SLX Optical Statistics

XCO supports optical statistics for SLX devices.

XCO collects the device inventory and channel media information during device discovery and stores it in the database. The device inventory information is refreshed every 15 minutes. XCO uses the `show media` and `show interface status` commands to construct the media optical information.

Related Topics

[View Device Inventory](#) on page 75

View Device Inventory

About This Task

Device details are displayed in cards by information type, which varies by the device you select. Different devices provide different information. Device details can include some or all of the following:

- Chassis, including type and serial number
- Line card, including name and up time
- Health, including system up-time and BIOS version
- Thermal, including sensor name and current temperature
- Fan, including status and speed
- PSU, including name and status
- LED, including name and state
- Media/optical levels, including TX Power and RX Power
- Port, including slot number and admin status

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.

3. Select the **Parts Details** tab.

The lists of details vary by device. The following is an example.

← MLXe_247_QA Device Actions

- Overview
- Policies and Configuration
- Parts Details**
- Monitor/Troubleshoot

Chassis Information	
Name	MLXe 8-slot
Last Updated	25 Aug, 2022 16:28
Type	MLXe 8-slot
Serial Number	BGB2522L005
Part Number	40-1000362-04
Firmware Revision	6.3.00bd

Line Card Information	
^ Linecard-1	
Name	Linecard-1
Last Updated	25 Aug, 2022 16:28
Type	BR-MLX-10Gx20 20-port 1/10GbE
Status	CARD_STATE_UP
Serial Number	-
Part Number	-

Fan Information	
^ Fan Unit-A-1	
Name	Fan Unit-A-1
Last Updated	25 Aug, 2022 16:28
Status	OK
Speed	LOW (50%)
v Fan Unit-A-2	
v Fan Unit-B-1	

PSU Information	
^ PS Unit-1	
Name	PS Unit-1
Last Updated	25 Aug, 2022 16:28
Type	-
Status	not present
v PS Unit-2	
v PS Unit-3	

Port Information	

2022 Extreme Networks Inc. [About](#) [Terms](#) [Privacy Policy](#) ExtremeCloud Orchestrator


Download Device Inventory

You can download the channel media information along with device inventory to a spreadsheet.

About This Task

To download multiple or bulk device inventory, see [Download Bulk Device Inventory](#) on page 33.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select the **Parts Details** tab.
4. In the upper left corner of the page, select  **Download Inventory**.
A file in .xlsx format is downloaded to your device.

Monitor/Troubleshoot (Packet Broker Mode)

Capturing Support Save logs is key to successful troubleshooting.

For more information, see [Support Save](#) on page 25.

Troubleshoot Configuration

Use the **Monitor/Troubleshoot** page to select a device configuration and view the statistics in the service chain.

About This Task

For any selected device, you can view the members of the related ingress and egress groups. You can also view the configuration of the related ingress policy, such as the protocol, the source IP address, and the Ethernet type.

Real-time statistics, such as packet flow and bit rate, can help you troubleshoot device issues. These statistics are available when you drill down to the Troubleshoot Configuration by Ingress Group page.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the **Devices** page, click anywhere in the required device row except the Actions column (•••) to proceed to the device Overview page.
3. Select the **Monitor/Troubleshoot** tab.

The Troubleshoot Configuration by Ingress Group page displays a list of ingress groups and their related ingress policies and egress groups.

The screenshot shows the 'Troubleshoot Configuration By Ingress Group' page in the ExtremeCloud Orchestrator GUI. The page displays a table with the following data:

Ingress Group	Ingress Policy	Egress Group	Actions
VNI-decap-NT	routemap_4 v6 v4	Tool-4 Tool-1 Tool-2 Tool-3	↗
Input-from-tap	Process-Access-list v4 v6 L2	Egress-to-Analytics-Tool	↗


The list of ingress policies shows the related protocol and layer.

4. To view details of an ingress group, such as its members, select the name of the group.

The Details pane is displayed at the bottom of the page.

The screenshot shows the details pane for the 'sintex_rt_ig' ingress group. The table displays the following information:


Members	Outer Tunnel Type	Ingress Policy
ethernet 4/2	-	sintex_rt

You can refresh the pane, filter for specific content, and add or remove columns. Use the filtering functions at the upper right of the pane. 

- To view details of a policy, such as its sequence number, select the name of the policy.

Use the filtering functions at the upper right of the Details pane to change the contents of the pane.

sintex_rt.ig > sintex_rt


Showing 1 - 3 of 3 results Page Size 10 Latest as on 10:01:10 AM 

Seq.#	IPv4 Match	IPv6 Match	L2 Match	UDA Match	Deny	Egress Group
56	-	-	sintex_l2	-	-	101 - vlan
57	-	sintex_v6	-	-	true	101 - vlan
65535	sintex_v4	-	-	-	-	101 - vlan


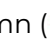
- To view the details of a egress group, such as its egresses, select the name of the group.

Use the filtering functions at the upper right of the Details pane to change the contents of the pane.





sintex_rt.ig > sintex_rt > 101 - vlan

Showing 1 - 2 of 2 results Page Size 10 Latest as on 10:01:46 AM 

Egresses	Members	Egress Policy
egress_vlan_101_e_2_4	ethernet 2/4	-
egress_vlan_101_e_4_5	ethernet 4/5	-

- To view the configuration of an ingress group, select the Troubleshoot () icon for the group from the Actions column ().

The Troubleshoot Configuration by Ingress Group page opens.

-  Overview
-  Policies and Configuration
-  Parts Details
-  Monitor/Troubleshoot

← Troubleshoot Configuration By Ingress Group

Ingress Interfaces:

ethernet 1/6:1


Policy Matches:


1_forward-any: #1 1_decide-filter-forward: #1

1_fwd-any: #1

Egress Interfaces:

7_Output-to-Tool1: #ethernet 2/2:3

 View Statistics

 Clear Stats

Ingress Group

▲ Input-from-tap 1 Ports/Port Channels

☑ Members

☑ ethernet 1/6:1

Ingress Policy

▲ Process-Access-list 1 ipv4 | 1 ipv6 | 1 I2

▲ 1 1 ipv4 | 1 ipv6 | 1 I2

The page displays expandable lists of ingress groups, ingress policies, and egress groups.

- a. To display configuration details, select one or more items in the expandable lists

The details are displayed at the top of the page. In the image, you can see selected interfaces and a matching policy.


- b. To view real-time statistics, select one or more configuration items and then select **View Statistics**.

A new page opens to display 3 panels of statistics, the contents of which vary depending on the configuration items you chose.

Troubleshoot Configuration By Ingress Group ×

Ingress Group Reset Counters					Ingress Policy Reset Counters				Egress Group Reset Counters				
PolicyUDA_Jg					PolicyUDA				Egress Group - 10				
Name	Octets	Broadca...	Multicas...	Packets	Match	Type	Bits Rate	Packets ...	Name	Octets	Broadca...	Multicas...	Packets
1/5	0/0	0/0	0/0	0/0	xyz	uda	-	-	1/20	0/0	0/0	0/0	0/0
1/6	55651640/51	0/0	412237/412	412237/412	xyz	uda	2160	2					

You can select **Reset Counters** to refresh the statistics.

You can add or remove columns and you can switch to a chart format. Use the functions at the upper right of each panel. 

- c. To clear statistics selections, select **Clear Stats**.

Firmware Upgrade

You can download and upgrade the firmware on multiple devices.

For information about deploying XCO, see the [ExtremeCloud Orchestrator Deployment Guide, 3.2.0](#).

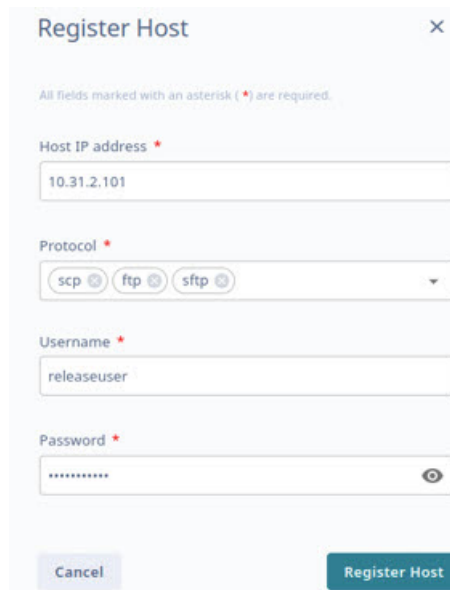
Register Firmware Host

Procedure

1. In the Navigation menu, select **Device Inventory**.

The **Devices** window opens.

2. Select **Settings > Firmware Hosts > Register Host**.



The **Register Host** window opens.

3. In the **Host IP address** field, provide the IPv4 or IPv6 address of the firmware host sever.
If a firmware host server has both IPv4 and IPv6 addresses, each IP address is treated as an independent entry.
4. In the **Protocol** field, select one or more options from the available protocols.
 - Packet Broker Mode:
 - SCP
 - TFTP
 - Fabric Mode:
 - FTP
 - SFTP
 - SCP
5. In the **Username** field, provide a name.
6. In the **Password** field, provide the password.
7. Select **Register Host**.

View Registered Firmware Hosts

Procedure

1. On the Navigation menu, select **Device Inventory**.
The **Devices** window opens.

2. Select **Settings > Firmware Hosts**.



The list of registered hosts opens.

Change a Firmware Host

Procedure

1. In the Navigation menu, select **Device Inventory > Settings > Firmware Hosts**.
The list of registered firmware hosts opens.
2. Select **Edit** from the Actions column (•••) for the firmware host IP address you want to edit.
3. Complete the fields as described in [Register Firmware Host](#) on page 79.

Delete a Firmware Host

Procedure

1. In the Navigation menu, select **Device Inventory > Settings > Firmware Hosts**.
The list of registered firmware hosts is displayed.
2. Select **Delete** from the Actions column (•••) for the host IP address you want to delete.

Upgrade Firmware (Device Level)

Before You Begin

- Register firmware host. For more information, see [Register Firmware Host](#) on page 79.
- When you upgrade to a new firmware image on SLX and Extreme 9920 products, the previous image is moved to the secondary location, and the previous secondary image is moved to the temporary location until the new image is committed.

Extreme 9920 devices overwrite the firmware in the specified location and may not have a secondary image available.

- If Extreme 9920 microservices do not come up within 25 minutes after the firmware upgrade, the image rollbacks automatically.
- If any of the Extreme 9920 microservices do not come up after image rollback, the device is set to **Degraded** state.

About This Task

For SLX devices, XCO extracts the target firmware version file name from the directory name.

Example:

```
/root/slxos18s.1.03/slxos18s.1.03a
Target firmware version: 18s.1.03a
```

Example:

Absolute path to the binary file for Extreme 9920 devices:

```
/root/TierraOS-21.1.2.0-NPB.bin
```

For MLX devices, the target firmware version file name is extracted from the manifest file name.

Example:

```
XMR-MLX/MLX06300bc_Manifest.txt
Target firmware version 6.3.00bc
```



Note

As a best practice, do not change the target firmware version file name and the directory name.

Table 13: Supported protocol

Device Type	Protocol
Extreme 9920	SCP
SLX Network Packet Broker (NPB)	SCP
SLX fabric	SCP, SFTP, FTP
MLX	TFTP

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the Devices page, select **Upgrade Firmware** from the Actions column (☰) for the device you want to upgrade.
Alternatively, click anywhere in the device row except the Actions column to proceed to the Device Overview page and select **Upgrade Firmware** from the Device Actions menu.
3. In the **Host** field, provide the IPv4 or IPv6 address of the firmware host server.
4. In the **Absolute Path** field, provide the firmware file path.
5. Select **Download and Activate**.
6. Select **Upgrade Firmware**.

Related Topics

[Register Firmware Host](#) on page 79

[View Registered Firmware Hosts](#) on page 80

[Change a Firmware Host](#) on page 81

[Delete a Firmware Host](#) on page 81

[Rollback Upgrade \(Packet Broker Mode\)](#) on page 83


Rollback Upgrade (Packet Broker Mode)

About This Task

XCO does not support firmware rollback for SLX and MLX devices.

Procedure

1. In the Navigation menu, select **Device Inventory**.
2. In the Devices page, select **Rollback Upgrade** from the Actions column (☰) for the device you want to roll back to the previous version.



Users

[Role Based Access Control](#) on page 84

[User Roles](#) on page 85

[Authentication Tokens](#) on page 85

[Local](#) on page 86

[Host](#) on page 88

[User Settings](#) on page 89

[Change Password](#) on page 94

[Logout](#) on page 95

Use the XCO user interface to configure the preferred authentication method for validating users.

XCO supports the following methods to manage and authenticate users:

- External LDAP server
- External TACACS+ server
- Local DB user
- Unix authentication on the host where XCO is installed

XCO supports predefined role management for LDAP and TACACS+. You can map the LDAP and TACACS+ specific roles with the predefined XCO roles.

For more information, see:

- [LDAP Settings](#) on page 90
- [TACACS+ Settings](#) on page 92

Role Based Access Control

XCO supports Role Based Access Control (RBAC). RBAC defines the capabilities that a user account has based on the assigned role. A role defines the access privileges of the user accounts.

XCO validates user privileges based on the assigned role:

- Custom roles are not supported. For information on supported roles, see [User Roles](#) on page 85.
- User-defined role management is supported for LDAP and TACACS+. For more information, see [LDAP Settings](#) on page 90 and [TACACS+ Settings](#) on page 92.

User Roles

A user is associated with one role. The user name and role of the logged-in user are displayed in the title bar.

Table 14: User role definitions

Role	Functions
SystemAdmin	Users with this role have complete privileges to perform all operations in the system. Note: The default host user who installs the XCO application has this role. You cannot edit or delete the host user.
NetworkOperator	Local users with this role have read-only privileges to all operations in the system. These users can change their own account password.
Fabric Mode Only:	
FabricAdmin	Users with this role have privileges to perform fabric management, device management, and location management operations.
TenantAdmin	Users with this role have read-only privileges to all operations in the system.
SecurityAdmin	Users with this role have privileges to perform user management operations.
SystemDebugger	Users with this role have privileges to perform system debug operations.

Authentication Tokens

Authentication tokens that are generated when a user logs in to XCO are stored in memory and validated for token authentication and authorization.

The token is cleared under the following conditions:

- User role modification
- User deletion
- User blocking
- User logout
- Session expiration
- Token expiration

If a user token is cleared during an active user session, the user is prompted to log in again.

Local

You can use the **Local** page to create and manage local users.

Add User

Only a user with the SystemAdmin role can add a local user.

About This Task

When the first local user is added, XCO automatically adds the **LOCAL Auth** type to the authentication preference settings in the following situations:

- **LOCAL auth** preference does not exist
- Authentication preference settings limit of five entries is not exceeded

Procedure

1. In the Navigation menu, select **Users**.
 2. Select **+ Add User**.
 3. In the **User Name** field, enter the user's user name.
 4. In the **User Role** field, select the required user roles.
 - **NetworkOperator**
 - **SystemAdmin**
 - Fabric mode only:
 - **FabricAdmin**
 - **SecurityAdmin**
 - **SystemDebugger**
 - **TenantAdmin** (created dynamically per tenant)
- XCO supports multiple role mapping for all users. For more information, see [User Roles](#) on page 85.
5. In the **New Password** and **Confirm New Password** fields, enter the new password for the user.
 6. In the **Email-id** field, enter the user's email address.

Special characters specified by RFC-5322 are supported in the email field.
 7. (Optional) Complete the other fields as required.
 8. Select **Save**.

Edit User


Before You Begin

Only a user with the role of SystemAdmin can change the role of another local user.

About This Task

To change the role of an LDAP or TACACS+ user, change the role on the remote server using the appropriate method.

Procedure


1. In the Navigation menu, select **Users > LOCAL**.
2. Select  for the relevant user.
3. Select **Edit User**.
4. In the **User Type** field, select **NetworkOperator** or **SystemAdmin**.
For more information, see [User Roles](#) on page 85.
5. Save your changes.

Block User

Before You Begin

Only a user with the SystemAdmin role can block or unblock a local user.

Procedure


1. In the Navigation menu, select **Users**.
2. Select the **Local** tab.
3. Select  for the relevant user.
4. Select **Block User** to block the user.

Unblock User

Before You Begin

Only a user with the SystemAdmin role can block or unblock a local user.

Procedure


1. In the Navigation menu, select **Users**.
2. Select the **Local** tab.
3. Select  for the blocked user.
4. Select **Unblock User** to unblock the user.

Request Reset Password

Before You Begin

- Only a user with the SystemAdmin role can reset the password of local users.
- Automated mail service for sharing the user password is not available.
- Password complexity check is not available.
- Local user passwords do not expire.

Procedure

1. In the Navigation menu, select **Users**.
2. Select  for the relevant user.
3. Select **Reset Password**.
The **Password Reset** window opens.

4. Enter the new password for the user.
5. Confirm the password.
6. Select **Save**.

The user is prompted to change the password on first login after password reset.

Change Password on First Login

You are prompted to change the password on first login.

Procedure

1. In the **New Password** field, enter the password.
2. In the **Confirm Password** field, enter the password again.
3. Select **Change Password**.

The password is changed and you are logged out of the user interface.

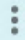
What to Do Next

Log in to the user interface using the new password.

Delete User

Only a user with the role of SystemAdmin can delete a local user.

Procedure

1. In the Navigation menu, select **Users > LOCAL**.
2. Select  for the relevant user.
3. Select **Delete User**.

Host

When XCO is deployed, the user who installs the application is configured as SystemAdmin with complete access and permissions.

Host user authentication is configured as the default authentication method.

Change Host User Role

The default host user who installs ExtremeCloud Orchestrator is automatically added to the host users role mapping page. You cannot edit or delete the default host user.

Procedure

1. In the Navigation menu, select **Users**.
2. Select **+ Add Host**.
3. From the **User Name** drop-down menu, select the required host user.
4. In the **User Type** drop-down menu, select the required user type:
 - **NetworkOperator**

- **SystemAdmin**
 - Fabric mode only:
 - **FabricAdmin**
 - **SecurityAdmin**
 - **SystemDebugger**
 - **TenantAdmin** (created dynamically per tenant)
5. Select **Save**.

User Settings

The **User Settings** page in the XCO user interface allows you to configure the LDAP and TACACS+ authentication settings and change the authentication level priority for the available authentication methods.

In the Navigation menu, select **Users > Settings** to access the **User Settings** page. You can access **User Settings** from all pages on User Management.

For more information, see [Authentication Settings](#) on page 89.

Authentication Settings

You can change the user authentication level priority among TACACS+, LDAP, Local, and HOST servers.

About This Task

XCO supports a maximum of five authentication settings. Host is the default authentication setting.

If a configured authentication level priority server is unreachable, failover to the next server occurs. If all servers decline to authenticate the user, the other configured authentication methods are attempted in the authentication order, and eventually the user is denied.

Procedure

1. In the Navigation menu, select **Users > Settings**.
2. In the **Authentication Settings** page, drag and drop the required server settings to reorder the authentication level priority.
3. Select **Apply** to save the changes.
4. To add a server to the existing authentication settings, select **Add Server**.
5. Select an authentication levels and then select **Apply**.

LDAP Settings

XCO supports Lightweight Directory Access Protocol (LDAP). The **Settings** page in the user interface allows viewing and managing of LDAP server configurations.

LDAP is an open-source protocol used for centralized authentication through directory service. If the configured LDAP servers decline to authorize the user, the other authentication methods are attempted in the order they are configured.

Active Directory (AD) is a directory service that supports a number of standardized protocols such as LDAP, Kerberos authentication, and Domain Name Server (DNS), to provide various network services. AD uses a structured data store as the basis for a logical, hierarchical organization of directory information. AD includes user profiles and groups as part of directory information, so it can be used as a centralized database for authenticating third-party resources.

XCO supports the following LDAP methods to authenticate users:

- The user role is included directly as an attribute in the user definition entry.
- The user has the `memberOf` entry or any appropriate group definition attribute to identify the groups assigned.
- The user entries are present in the LDAP group definition.
- If the user entry is not present or not mapped to the correct predefined role in XCO, the user login fails. For more information, see [Map an LDAP User Role](#) on page 92.



Note

If LDAP group definition methods are used for user authentication, the corresponding LDAP group must be mapped to an user role in XCO.

- XCO supports a maximum of five LDAP servers.
- The LDAP configuration name must be unique for configuring the authentication policy.

Create/Update LDAP Server

You can add LDAP connection details so that LDAP users can sign in to the XCO user interface.

About This Task

When a new LDAP server is added, XCO automatically adds it to the authentication preference settings if the authentication preference limit of five entries is not exceeded.

Procedure

1. In the Navigation menu, select **Users**.
2. Select **Settings > LDAP Settings > Add LDAP Server**.
Alternatively, you can select **LDAP > Connect LDAP** to configure the first LDAP server.
3. In the **Name** field, enter a name for the LDAP server.
The name can contain up to 32 alphanumeric characters without spaces.

4. (Optional) If multiple LDAP servers are available, proceed to the next step. Else, go to step 6.
5. In the **Host** field, enter the host name, IPv4, or IPv6 address of the LDAP server.
6. (Optional) In the **Port** field, enter the TCP port used for authentication.
7. (Optional) In the **CA Certificate** field, enter the CA certificate location.
Select the CA certificate to use when validating the server certificate that the LDAP server sends. The CA certificate must be issued by the same CA that issued and signed the server certificate for the LDAP server.
8. In the **Timeout(Secs)** field, enter the timeout value in seconds.
The default timeout value is 5 seconds.
9. (Optional) In the **Bind User Name** field, enter the LDAP server user name.
The Bind User Name is used for authenticating the LDAP server when initiating a connection.
10. (Optional) In the **Bind User Password** field, enter the password for the LDAP server.
The Bind User Password is used for authenticating the LDAP server when initiating a connection.
11. In the Advanced section, complete the following fields as required:
 - **User Search Base:** Specifies the name of the node from which to start searching for users.
 - (Optional) **User Object Class:** Specifies the name of the user object class. The default value is `inetOrgPerson`.
 - (Optional) **User Login Attribute:** Specifies the login username attribute. The default value is `uid`.
 - (Optional) **User Role Attribute:** Specifies the user role attribute.
 - (Optional) **User Role Attribute Key:** Specifies key to the user role attribute.
 - (Optional) **User Member Attribute:** Specifies the member attribute of the user.
 - (Optional) **Group Search Base:** Specifies the name of the node from which to start searching for groups.
 - (Optional) **Group Object Class:** Specifies the name of the group object class. The default value is `groupOfNames`.
 - (Optional) **Group Attribute:** Specifies the group attribute. The default value is `cn`.
 - (Optional) **Group Member User Attribute:** Specifies the group member user attribute. The default value is `entrydn`.
 - (Optional) **Group Member Mapping Attribute:** Specifies the group member mapping attribute. The default value is `member`.
 - (Optional) **TLS check box:** Enables LDAP over SSL/TLS
 - (Optional) **Insecure-TLS check box:** Enables LDAP without certificate verification
12. Select **Test Connection and Save** to save your selections.
The **Authentication Settings** page displays the new configuration.

What to Do Next

[Map an LDAP User Role](#) on page 92

Map an LDAP User Role

You can map a local LDAP role to one of the pre-defined XCO roles.

About This Task

The LDAP server name is used as `auth-identifier` for mapping a LDAP user role.

Procedure

1. In the Navigation menu, select **Users**.
2. Select **Settings > LDAP Settings**.
3. Select an LDAP server.
4. Select **Map User Roles**.
5. Select the required **LDAP User Roles**.
 - **GROUP**
 - **USER**
6. Select the required **XCO User Roles** to map the local LDAP role.
7. Select **Apply**.

The LDAP server page displays the new mapping.

TACACS+ Settings

Terminal Access Controller Access-Control System Plus (TACACS+) is an external authentication server used for verifying user credentials.

The TACACS+ protocols support environments that are configured for authentication, authorization, and accounting (AAA) services. When TACACS+ is configured through the XCO interface, TACACS+ users can log in to the XCO interface.

XCO supports TACACS+ authentication in the following ways.

- XCO supports up to four TACACS+ servers for authentication purposes and contacts them in the order they are configured.
- The user roles specified in the TACACS+ server configuration can be one of the following.
 - One of the supported XCO roles: NetworkOperator and SystemAdmin. For more information, see [User Roles](#) on page 85.
 - A local TACACS+ role that you can map to XCO. For more information, see [Map a TACACS+ User Role](#) on page 93.
 - The `xco-role` attribute must be included in the TACACS+ configuration file.
 - If the `xco-role` attribute is not present or not mapped to the correct predefined role in ExtremeCloud Orchestrator, the user login fails.
- TACACS+ authentication must be enabled. If TACACS+ authentication is not enabled, only local authentication is used.
- If remote authentication fails, XCO attempts to use local authentication, which is successful only if the user is in the XCO database.

- The secret key configured for XCO must be the same as the secret key from the TACACS+ server configuration file. Authentication fails if the two values do not match.
- XCO supports two TACACS+ authentication protocols.
 - **CHAP**: Challenge Handshake Authentication Protocol
 - **PAP**: Password Authentication Protocol

Add TACACS+ Server

You can add TACACS+ connection details so that TACACS+ users can sign in to the XCO interface.

About This Task

When a new TACACS+ server is added, XCO automatically adds it to the authentication preference settings if the authentication preference limit of five entries is not exceeded.

Procedure

1. In the Navigation menu, select **Users**.
2. Select **Settings > TACACS+ Settings > Add TACACS+ Server**.
Alternatively, you can select **TACACS+ > Connect TACACS+** to configure the first TACACS+ server.
3. In the **Host** field, enter the IPv4 or IPv6 address of the TACACS+ server, in CIDR format.
4. In the **Port** field, enter the TCP port used for authentication.
The default authentication port is 49.
5. In the **Secret Key** field, enter the shared secret that enables messages between the client and the TACACS+ server.
The value you enter must match the shared secret in the TACACS+ server configuration file.
6. In the **Protocol** field, select one of the following authentication protocols.
 - **CHAP**: Challenge Handshake Authentication Protocol
 - **PAP**: Password Authentication Protocol
7. Select **Test Connection and Save** to save your selections.
The Settings page displays the new configuration.

What to Do Next

[Map a TACACS+ User Role](#) on page 93

Map a TACACS+ User Role

You can map a local TACACS+ role to one of the pre-defined XCO roles.

About This Task

The TACACS+ server `host` is used as `auth-identifier` for mapping a TACACS+ user role.

Procedure

1. In the Navigation menu, select **Users**.
2. Select **Settings > TACACS+ Settings**.
3. Select a TACACS+ server.
4. Select **Map User Roles**.
5. Select the required **TACACS+ User Roles**.
6. Select the required **XCO User Roles** to map the local TACACS+ role.
7. Select **Apply**.

The TACACS+ server page displays the new mapping.

Change a Server Configuration

You can change the configuration of a LDAP or TACACS+ server for accessing the XCO interface.

Procedure

1. In the Navigation menu, select **Users**.
2. Select **Settings > LDAP Settings** or **Settings > TACACS+ Settings** as required.
Alternatively, you can select **LDAP** or **TACACS+** tab.
3. Select **Edit** for the server configuration that you want to change.
4. Update the server configuration as required.
5. Save your selections.
6. Select **Apply**.

The Authentication Settings page displays the changed configuration.

Delete a Server Configuration

You can delete the configured LDAP and TACACS+ host servers.

Procedure

1. In the Navigation menu, select **Users**.
2. Select the **TACACS+** or **LDAP** tab.
3. Select **Delete** for the server configuration that you want to delete.
Alternatively, you can delete the LDAP and TACACS+ server configurations from **Users > Settings > Authentication Settings** .

Change Password

Logged-in users can change their own passwords.

Procedure

1. From the User Profile menu, select **Change Password**.
2. Update the password.

Logout

Procedure

From the User Profile menu, select **Logout**.



Logs

[System Logs](#) on page 96

[User Logs](#) on page 97

The XCO user interface enables viewing of System logs and User logs. The System logs persist for two hours, and User logs persist for a week.

Exporting System logs and User logs is not supported in XCO.




System Logs

System logs describe the status of monitored devices.

About This Task

System logs are based on RASLog notifications. The system logs are stored for a duration of two hours.

Procedure

1. In the Navigation menu, select **Logs > System**.
The system logs provide the following information:
 - IP address
 - Hostname
 - Severity
 - Message
 - Date
2. Use the **Search** bar to look up a system log.
3. Use the refresh icon  to reload the system logs.
4. Use the filter option  to view the system logs by **Severity** or **Message**.
 - a. In the **System Logs** widget, select .
 - b. Select a system log value between **Severity** or **Message**, and then enter the filter value.
 - c. Select **Add Filter** to include more filter options, or **Apply Filter** to view the system logs based on your previous selection.

User Logs

You can view user logs to understand the transactions that a user has performed.

About This Task

XCO offers several types of logs related to user transactions: Device, Device Config, and User Login. These logs provide the following information.

Procedure

1. In the Navigation menu, select **Logs > User**.
2. To view user transactions on devices, select **Device**.

The device logs provide the following information:

Table 15: User device logs

Log Type	Description
Device	Device add or delete transactions: <ul style="list-style-type: none"> • User name • Action, such as delete or discover a device • IP address • Location • Status, such as success or failed • Error message to explain a failure • Date

3. To view user transactions related to configuration, select **Device Config**.

The device config logs provide the following information:

Table 16: User device config logs

Log Type	Description
Device Config	Device configuration transactions: <ul style="list-style-type: none"> • User name • Action, such as add, update, clearing counters, packet capture, or delete a configuration • IP address • Location • Status, such as success or failed • Error message to explain a failure • Date

4. To view user transactions related to logging in, select **User Login**.

The user login logs provide the following information:

Table 17: User login logs

Log Type	Description
User Login	User login and logout transactions: <ul style="list-style-type: none">• User name• Action, such as log in or log out• User role• Log in time• Whether the action was successful



Library (Packet Broker Mode)

[Matches](#) on page 99

[Policies](#) on page 103

[UDA](#) on page 105

The **Library** provides access to policy rule matches, policies, and user-defined ACL (UDA) profiles (for SLX and MLX devices only).

Use the **Library** page to create, edit, export, clone, or delete policy rule matches, policies, and UDA profiles.

Matches

A policy rule match identifies the parts of a packet header that a rule targets, such as the source port or the payload length. On the Matches page, you can see matches and their associated device type, rule type, and number of rules. The page provides access to creating, exporting, cloning, editing, and deleting match-related functions.

Create a Policy Rule Match in the Library

Policy rule matches in the library can be imported to a device.

About This Task

To create a policy rule match for a specific device, see [Create a Policy Rule Match for a Device](#) on page 48.

Procedure

1. In the Navigation menu, select **Library > Matches > Add Match**.
2. In the **Name** field, enter a unique name for the match.
 - Alphanumeric characters, dashes, and underscores are allowed in the **Name** field.
 - The name, `a11` is a reserved keyword on 9920 and cannot be used.
3. In the **Device Type** field, select the required device type.
 - **9900**
 - **MLX**
 - **SLX**
4. In the **Type** field, select whether the match applies to IPv4, IPv6, L2, or UDA.
If you selected UDA on an SLX device, proceed to the next step. Else, go to step 7.

5. In the **Sub Type** field, select the appropriate match.
 - **Standard**: Matches the source address information
 - **Extended**: Matches the source and destination address information
6. In the **UDA** field, select a profile.
7. In the Match section, complete the following fields to identify the packets of interest.

**Note**

All fields are not mandatory. You can leave the fields blank unless noted.

The items that you can select vary by your selection in the **Protocol** field. The following describes all possible selections.

- **Protocol**: The protocol that you want to target. If the protocol you want is not in the list, select **None** and provide the ID of the protocol you want in the **Protocol ID** field. Every protocol has a numeric value that is defined by IETF.
- **Sequence**: The order in which this rule is performed in the match.
- **Protocol ID**: The ID of a protocol that you want to target. Use only when the protocol you want is not available in the **Protocol** field.
- **Source IP**: The IPv4 or IPv6 address of the device that sends the packets.
- **Source Mask**: The mask for the source IP address, in the following format: 255.255.255.255.
- **Destination IP**: The IPv4 or IPv6 address of the device that is to receive the packets.
- **Destination Mask**: The mask for the destination IP address, in the following format: 255.255.255.255.
- **Source Mac**: The MAC address of the device that sends the packets, in the following format: 1111.1111.1111 or 11:11:11:11:11:11. Any alpha characters in the address must be lowercase.
- **Source Mac Mask**: The mask for the source MAC address, in the following format: ffff.ffff.ffff or ff:ff:ff:ff:ff:ff. Any alpha characters in the mask must be lowercase.
- **Destination Mac**: The MAC address of the device that is to receive the packets, in the following format: 1111.1111.1111 or 11:11:11:11:11:11. Any alpha characters in the address must be lowercase.
- **Destination Mac Mask**: The mask for the destination MAC address, in the following format: ffff.ffff.ffff or ff:ff:ff:ff:ff:ff. Any alpha characters in the mask must be lowercase.
- **Source Port**: The port through which packets enter the device.
- **Source Port End**: The last port in the range of ports through which packets enter the device.
- **Destination Port**: The port through which packets leave the device. Valid values range from 1 through 65535.
- **Destination Port End**: The last port in the range of ports through which packets leave the device. Valid values range from 1 through 65535.
- **IP Payload Length**: The length of the IP packets that you want to target, or the size of the IP payload. Valid values range from 64 through 9000.

- **IP Payload Length End:** The last acceptable value of the IP payload. Valid values range from 65 through 9000.
- **DSCP:** The value of the Differentiated Services Code Point in the Type of Service field in the header. Valid values range from 0 through 63.
- **VLAN:** The VLAN ID. Valid values range from 0 through 4095.
- **EtherType:** Identifies the protocol that is encapsulated in the payload. For example, the EtherType value for IPv4 is 0x0800. Valid values range from 1536 through 65536 (numerical), or 0x0600 through 0xffff (hexadecimal), or are one of the following: ARP, IPv4, or IPv6.
- **PCP:** The Priority Code Point, a 3-bit field in a VLAN header. Valid values range from 0 through 7.
- **Tunnel ID:** The ID number of the tunnel. Valid values range from 1 through 16777215.
- **MATCH0, MATCH1, MATCH2, MATCH3:** Specifies the UDA Hexadecimal. SLX presents these as specific header fields such as `NEXT_HEADER`.



Note

- MLX UDA requires a match and mask for all fields.
- Use a mask of all zeros to make the any value for a field.

- **MASK0, MASK1, MASK2, MASK3:** Specifies the UDA Hexadecimal value used to mask the MATCH values. Use 0 bits for any value. A bit value of 1 must be matched.

8. In the Fragmentation section, select one or more of the following.

The items in this section vary by your selection in the **Type**, **Sub Type** and **Protocol** fields. The following list describes all possible selections.

- **Fragmented:** Targets target fragmented packets.
- **Non Fragmented:** Targets non-fragmented packets.
- **None:** Targets packets in which the DF (Don't Fragment) flag is set in the IP header.

9. In the Options sub-section, select one or more of the following:

The items in this section vary by your selection in the **Type**, **Sub Type** and **Protocol** fields, in particular selection of a Layer4 protocol such as UDP, TCP, or STCP. The following list describes all possible selections.

- **Acknowledgment:** Targets packets in which the ACK flag is set in the TCP header.
- **Congestion:** Targets packets in which the CWR flag is set in the TCP header.
- **ECN-Echo:** Targets packets in which the ECE flag is set in the TCP header.
- **Last Packet:** Targets packets in which the FIN flag is set in the TCP header.
- **Push:** Targets packets in which the PSH flag is set in the TCP header.
- **Reset:** Targets packets in which the RST flag is set in the TCP header.
- **Synchronize:** Targets packets in which the SYN flag is set in the TCP header.
- **Urgent:** Targets packets in which the URG flag is set in the TCP header.

10. In the Action section, select one or more actions to perform on the targeted items.
The items in this section vary by your selection in the **Protocol** field. The following list describes all possible selections.
 - **Drop** to drop the packet
 - **Count** to keep track of the number of packets that match the policy rule
 - **Log** to add the transaction to the log.
11. Select **Add**.
The match parameters (the new rule) appear in the pane on the right.
12. Repeat steps 7 through 11 until you have added all the rules you need.
13. Select **Save**.



Change a Policy Rule Match in the Library

You can add, change, or delete one or more rules in a policy rule match.

About This Task

To change a policy rule match for a specific device, see [Create a Policy Rule Match for a Device](#) on page 48.

Procedure

1. In the Navigation menu, select **Library > Matches**.
2. Select **Edit** () from the Actions column () for the policy rule match that you want to change.
3. Follow the instructions in [Create a Policy Rule Match in the Library](#) on page 99 to add, change, or remove rules from the match.
A new match is created with the updated configuration.

Export a Policy Rule Match from the Library

From the library, you can export a policy rule match to selected devices.

About This Task


You can export a rule match created either in the library or reconciled from another device to a set of devices.



Note

The rule matches can be exported to devices of the matching device type only. For example: 9920 to 9920 or SLX to SLX.

Procedure

1. In the Navigation menu, select **Library > Matches**.
2. In the Matches page, select match that you want to export.
3. In the Actions column, select **Export** ().
4. Select the devices to which you want to export the selected match.

5. Select **Export**.

Clone a Policy Rule Match

From the library, you can clone (copy) a policy rule match to create a new match with the same or similar configuration.

Procedure

1. In the Navigation menu, select **Library > Matches**.
2. In the Actions column, select **Clone** for the match that you want to copy.
3. In the **Name** field, provide a new name for the cloned match.
4. Save your selections.

Delete a Policy Rule Match from the Library

You can delete a policy rule match from the library.



Before You Begin

You cannot delete a match that is attached to any device.

About This Task

To delete a policy rule match from a specific device, see [Delete a Policy Rule Match from a Device](#) on page 52.

Procedure

1. In the Navigation menu, select **Library > Matches**.
2. Select one or more matches to delete.
3. Select **Delete** () from the Actions column () for the match you want to delete.

Policies

Ingress and egress policies define the actions to apply to inbound and outbound packets. On the Policies page, you can see policies and their associated device type, policy type, and number of rules. The page provides access to policy-related functions such as creating, exporting, cloning and deleting. For more information, see [Policies](#) on page 44.

Create a Policy in the Library

About This Task

Policies in the library can be exported to one or more devices. For more information, see [Export a Policy](#) on page 104.

Procedure

1. In the Navigation menu, select **Library > Policies > Add Policy**.
2. Follow the instructions in [Create an Egress Policy for a Device](#) on page 44 or [Create an Ingress Policy for a Device](#) on page 46.

Change a Policy in the Library

You can add, change, or delete one or more rules or actions in a policy.

About This Task

You can change a policy for a specific device or change a policy in the library.

Procedure

1. In the Navigation menu, select **Library > Policies**.
2. In the Policies page, select the policy that you want to change.
3. In the Actions column, select **Edit**.
4. Follow the instructions in [Create an Ingress Policy for a Device](#) on page 46 or [Create an Egress Policy for a Device](#) on page 44 to add, change, or remove rules or actions in the policy.



Tip

To reinstate a rule that is not included in the policy (the **Deny** field is selected), clear the **Deny** field.

Export a Policy

From the library, you can export a policy to selected devices.

About This Task


You can export a policy created either in the library or reconciled from another device to a set of devices.



Note

The policies can be exported to same device type devices only. For example: 9920 to 9920 or SLX to SLX.

Procedure

1. In the Navigation menu, select **Library > Policies**.
2. In the Policies page, select the policy that you want to export.
3. In the Actions column, select **Export** ().
4. Select the devices to which you want to export the selected policy.
5. Select **Export**.


Clone a Policy

From the library, you can clone (copy) a policy to create a new policy with the same or similar configuration.

About This Task

After cloning a policy, the policy can be exported to the same device type devices only. For more information, see [Export a Policy](#) on page 104.

Procedure

1. In the Navigation menu, select **Library > Policies**.
2. In the Library menu, select the policy that you want to copy.
3. In the Actions column, select **Clone** ().
4. In the **Name** field, provide a new name for the cloned policy.
5. Follow the instructions in [Create an Ingress Policy for a Device](#) on page 46 or [Create an Egress Policy for a Device](#) on page 44 to add, change, or remove rules from the policy.
6. Save your selections.

Delete a Policy in the Library

You can delete a policy from a device or from a library.

Before You Begin

- If a policy is associated with a device, it cannot be deleted from the library. Follow the instructions in [Delete a Policy from a Device](#) on page 48 to delete the policy from any associated device.
- Remove the policy from any associated ingress group. For more information, see [Change an Ingress Group](#) on page 57.

Procedure

1. In the Navigation menu, select **Library > Policies**.
2. Select one or more policies to delete.
3. Select **Delete**.

UDA

UDA is a set of rules defined to filter the network traffic. Packets are forwarded or dropped based on criteria specified in UDA.

XCO supports the following UDA configurations:

- UDA Match
- UDA Profile

A complete MLX or SLX UDA consists of a UDA profile and a UDA match.

In Extreme 9920, packets that regular ACLs cannot filter, require filtering based on deep packet inspection or a combination of MAC and IP fields. UDAs, also known as Flex ACLs provide deeper and flexible filtering.

Create an MLX UDA Profile in the Library

For MLX devices, you can create a UDA profile in the library or device inventory.

About This Task

To create an MLX UDA profile for a specific device, see [Create an MLX UDA Profile for a Device](#) on page 53.

Procedure

1. In the Navigation menu, select **Library > UDA > Add UDA Profile**.
2. In the **Name** field, enter a unique name for the UDA profile.
3. In the **Device Type** field, select **MLX**.
4. In the four **Offset** fields, select the appropriate offset values.
 - The offset value must be a multiple of 4 between 0 and 124.
 - An offset indicates the index of the received packet. For example, an offset of 0 indicates the first byte of the received packet.
5. Save your selections.

Create an SLX UDA Profile in the Library

A UDA profile can be associated with a UDA match.

About This Task

To create an SLX UDA profile for a specific device, see [Create an SLX UDA Profile for a Device](#) on page 53.

Procedure

1. In the Navigation menu, select **Library > UDA > Add UDA Profile**.
2. In the **Name** field, enter a unique name for the UDA profile.
3. In the **Device Type** field, select **SLX**.
4. Define the header fields that are required for a match.

The header fields you select constitute the header stack. As you select header types and header fields, additional header selections become available. The additional selections vary based on your header choices.

- a. In the **Header 0 Ethernet - Ethernet** row, select the field that is required for a match and then click **+** to add your selection.
 - b. In the **Header 1** row, select the type and field that are required for a match and then click **+** to add your selection.

Your selections determine whether a Header 2 row is displayed.
 - c. Make selections in the Header 2 row and in all subsequent rows until no more rows are available or until your header stack is complete.

A maximum of 4 Headers are supported in a UDA profile.
5. Save your selections.

Change a UDA Profile in the Library

You can change the parameters of UDA profile.

About This Task

To change a UDA profile for a specific device, see [Change a UDA Profile for a Device](#) on page 54.

Procedure

1. In the Navigation menu, select **Library > UDA**.
2. In the UDA Profile page, select **Edit** from the Actions column (⋮) for the UDA Profile that you want to change.
3. Follow the instructions for the type of UDA you are changing.
 - [Create an MLX UDA Profile in the Library](#) on page 106
 - [Create an SLX UDA Profile in the Library](#) on page 106

Delete a UDA Profile in the Library

You can delete a user-defined access list (UDA) profile from the library or device inventory page.

About This Task

To delete UDA profile from a specific device, see [Delete a UDA Profile from a Device](#) on page 54.

Before You Begin

You cannot delete a UDA profile that is attached to any device, match, or ingress-group.

Procedure

1. In the Navigation menu, select **Library > UDA**.
2. In the UDA Profile page, select one or more UDAs to delete.
3. Select **Delete Profile**.

Export a UDA Profile

You can clone (copy) a user-defined access list (UDA) profile to create a new profile with the same or similar configuration.

Procedure

1. In the Navigation menu, select **Library > UDA**.
2. In the UDA Profiles page, select **Export** (↑) from the Actions column (⋮) for the profile that you want to export.
3. In the Devices list, select the devices to which you want to export the configuration.
4. Select **Export**.

The configuration is exported to the destination devices.



Fabrics (Fabric Mode)

- [Create a Non-Clos Fabric](#) on page 109
- [Create a 3 Stage Clos Fabric](#) on page 113
- [Create a 5 Stage Fabric](#) on page 118
- [View Fabric Topology](#) on page 124
- [Download Fabric Inventory](#) on page 125
- [Delete Fabric](#) on page 126
- [Download Health Report](#) on page 126
- [Fabric Actions](#) on page 126

A fabric denotes a collection of interconnected devices in a topology on which underlay and overlay networks are configured.

XCO 3.2.0 supports building and managing small data center (non-Clos) fabrics and 3-Stage and 5-Stage IP Clos fabrics based on a BGP underlay with a BGP or EVPN overlay.

- Non-Clos topology involves one to four interconnected racks. Each rack consists of a pair of leaf nodes or border leaf nodes.
- 3-Stage Clos topology involves a spine layer and leaf or border leaf layer. The border leaf can be single-homed or dual-homed.
- 5-Stage Clos topology involves a super spine layer, spine layer, and leaf or border leaf layer. The leaf or border leaf can be single-homed or dual-homed.

For more information on IP fabric topologies, see [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).

You can use the **Fabrics** page in XCO to configure and manage IP fabrics.

For information on the ExtremeCloud Orchestrator user interface and common operations in the interface, see:

- [User Interface](#) on page 19
- [Refresh Page View](#) on page 20
- [Pagination](#) on page 21
- [Search, Group, and Filter](#) on page 21

Create a Non-Clos Fabric

Before You Begin

- A non-clos topology supports a maximum of four racks with two devices each.
- The devices must be registered with the inventory before adding them to the fabric.

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the Fabrics page, select **Create Fabric**.
3. In the **Fabric Name** field, enter a name for the fabric.
4. In the **Fabric Description** field, enter a description for the fabric.
5. Select **Non Clos** topology.


The screenshot shows a three-step wizard: 1. Non Clos, 2. Properties, and 3. Topology Validation. In the 'Non Clos' step, the user has entered 'NonCLOSMultirack' in the 'Fabric Name' field and 'NonCLOSMultirack' in the 'Fabric Description (Optional)' field. Three topology options are presented:

- Non Clos**: A diagram showing two racks, each with two leaf nodes (LF) connected in a mesh. Description: "Non CLOS topology involves n (1 to 4) number of racks interconnected to each other. Rack consists of a pair of leaf or a pair of border leaf".
- 3 Stage Clos**: A diagram showing two spine nodes (SP) connected to three leaf nodes (LF). Description: "3-Stage CLOS topology involves Spine Layer and Leaf/Border Leaf Layer. Leaf/Border Leaf can be single-homed or dual-homed.".
- 5 Stage Clos**: A diagram showing one super spine node (SSP) connected to two spine nodes (SP), which are each connected to two leaf nodes (LF). Description: "5-Stage CLOS topology involves Super Spine Layer, Spine Layer and Leaf/Border Leaf Layer. Leaf/Border Leaf can be single-homed or dual-homed.".

A 'Next' button is visible in the bottom right corner of the wizard.

6. Select **Next**.

- In the **Fabric Properties** page, complete the fields as required or select the **Auto Config Mode** check box to use the default fabric settings.



- To create a multi-rack Non-Clos fabric, clear the **Single Rack Deployment** check box.
- Select **Create Fabric**.
- In the **Physical Topology** page, add racks as required.

Use **Topology View** (🔗) and **List view** (☰) to switch the view between topology and list.

- Select **+** and **-** to add or remove a rack.



Alternatively, you can do the following:

- In the Devices panel, select **+ Add Rack** to add a new rack.
- From the rack menu (⋮), select **Remove** or select **Delete** (🗑️) in the devices panel to remove a rack.

- b. (Optional) From the rack menu (⋮), select **Convert to Border Rack** or **Convert to Leaf Rack** to change the rack type.

Physical Topology

10.20.246.6 added to the fabric successfully

Unassigned Devices

Name	Model	IP	Firmware	LastRef
SLX	3001	10.20.246.18	18s.1.03	10 minutes
SLX	3001	10.20.246.24	18s.1.03	3 minutes
Freedom-07	3001	10.20.246.23	18s.1.03	3 minutes
SLX	3001	10.20.246.17	18s.1.01c	3 minutes
NHF-Leaf1	3009	10.20.246.5	20.4.3sksos20.4.3a_230218_1918	10 minutes

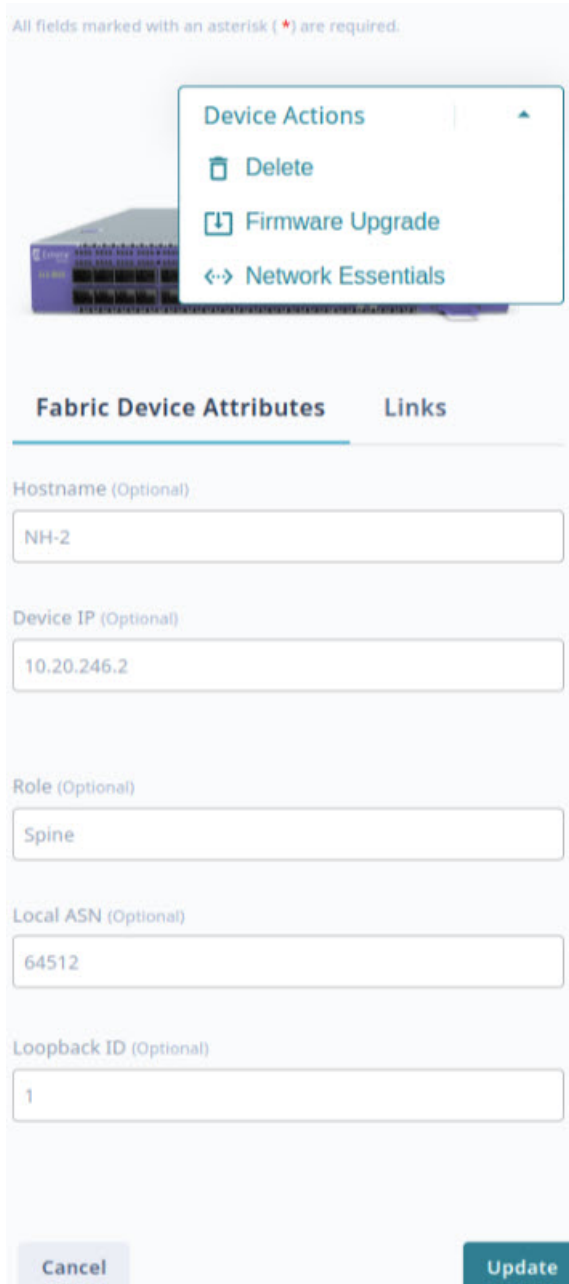
Showing 6 - 10 of 12 results

Assigned Devices 3/4

Name	Model	Type	Ip	Ports	Firmware
Rack1					
NH-1	3012	Leaf	10.20.246.1	-	20.4.3sksos20.4.3a_230218_1918
NH-2	3012	Leaf	10.20.246.2	-	20.4.3sksos20.4.3a_230218_1918

11. Drag and drop the required devices from the Devices panel to the rack.
- Select **Add Device** to add a device to the inventory. For more information, see [Add Devices](#) on page 31.
 - The devices available in the rack are displayed in the **Assigned Devices** list.
 - The inventory devices that are not part of the fabric are displayed in the **Unassigned Devices** list.
 - You can select devices in rack and update fabric specific configurations such as ASN, VTEP Loopback ID, and Loopback ID. The attributes vary by device role.
 - You can select and edit device and fabric configurations directly from the **Physical Topology** or **Devices** panel as required.
 - From the **Device Actions** menu, select **Network Essentials** to modify network essential configurations of the required device ports. For more information, see [Network Essentials](#) on page 129.

All fields marked with an asterisk (*) are required.



Device Actions

- Delete
- Firmware Upgrade
- Network Essentials

Fabric Device Attributes Links

Hostname (Optional)
NH-2

Device IP (Optional)
10.20.246.2

Role (Optional)
Spine

Local ASN (Optional)
64512

Loopback ID (Optional)
1

Cancel Update

12. Select **Discover Topology** to validate the topology.

The discovered topology is displayed.

The ports or links that are down are marked in red in the topology. To turn a port or link up:

- Select the link that is down.
- In the **Ports Up** dialog box, configure the ports.
- Select **Apply**.
- Select **Discover Topology** to validate the topology.

The screenshot displays the 'Physical Topology' view in the ExtremeCloud Orchestrator GUI. The main area shows a network diagram with two racks, Rack1 and Rack2. Rack1 contains two leaf nodes with IP addresses 10.20.246.1 and 10.20.246.2. Rack2 contains two leaf nodes with IP addresses 10.20.246.5 and 10.20.246.6. Green lines represent connections between the leaf nodes in Rack1 and Rack2. On the right side, there are two panels: 'Unassigned Devices' and 'Assigned Devices 4/4'. The 'Unassigned Devices' panel shows a table with columns for Name, Model, IP, Firmware, and LastRef. The 'Assigned Devices' panel shows a table with columns for Name, Model, Type, Ip, Ports, and Firmware.

Name	Model	IP	Firmware	LastRef
SLX	3001	10.20.246.18	18s.1.03	11 minutes
SLX	3001	10.20.246.24	18s.1.03	4 minutes
Freedom-07	3001	10.20.246.23	18s.1.03	4 minutes
SLX	3001	10.20.246.17	18s.1.01c	4 minutes
NH-Leaf2	3012	10.20.246.4	20.4.3slxos20.4.3_221117_0600	54 minutes

Name	Model	Type	Ip	Ports	Firmware
Rack1					
NH-1	3012	Leaf	10.20.246.1	0/21	20.4.3slxos20.4.3a_230218_1918
NH-2	3012	Leaf	10.20.246.2	0/21	20.4.3slxos20.4.3a_230218_1918
Rack2					

13. To enlarge the topology to the width of the interface, select **Expand** ().

14. To zoom in and out on the topology screen, use the **Zoom** ( ) icons.

15. To scroll through the topology screen, use the **Scroll** () icon.

16. Select **Finish** to configure the topology.

The non-Clos fabric topology is configured.

What to Do Next

Select **View Fabric** or **Proceed to Dashboard** to return to the **Fabrics** page.

Related Topics

[View Fabric Topology](#) on page 124

[Create a 3 Stage Clos Fabric](#) on page 113

[Create a 5 Stage Fabric](#) on page 118

Create a 3 Stage Clos Fabric

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the Fabrics page, select **Create Fabric**.
3. In the **Fabric Name** field, enter a name for the fabric.
4. In the **Fabric Description** field, enter a description for the fabric.

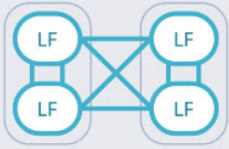
5. Select **3 Stage Clos** topology.

1 3 Stage Clos 2 Properties 3 Select Devices 4 Topology Validation

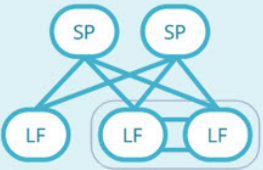
Add Fabric Name and Select Type

Fabric Name *

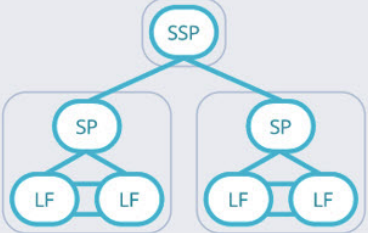
Fabric Description (Optional)



Non Clos
Non CLOS topology involves n (1 to 4) number of racks interconnected to each other. Rack consists of a pair of leaf or a pair of border leaf



3 Stage Clos
3-Stage CLOS topology involves Spine Layer and Leaf/Border Leaf Layer. Leaf/Border Leaf can be single-homed or dual-homed.




5 Stage Clos
5-Stage CLOS topology involves Super Spine Layer, Spine Layer and Leaf/Border Leaf Layer. Leaf/Border Leaf can be single-homed or dual-homed.

[Next](#)

6. Select **Next**.


7. In the **Fabric Properties** page, complete the fields as required or select the **Auto Config Mode** check box to use the default fabric settings.



8. Select **Create Fabric**.

9. Select the check boxes of the required leaf devices from the following tabs:

- **Spines**
- **Border Leafs**
- **Leafs**



The border leaf devices are optional. Select **Skip Border Leafs** to skip border leaf devices.

Select **+ Add Device** to add new devices to the inventory. For more information, see [Add Devices](#) on page 31.

You can select any device row and update fabric specific configurations such as ASN, VTEP Loopback ID, and Loopback ID. The attributes vary by device role.

10. Select **Accept Spine Leafs**, **Accept Border Leafs**, or **Accept Leafs** as applicable.

11. Select **Accept All** to add all devices to the topology.

The screenshot displays the GUI interface for fabric configuration. At the top, a progress bar shows four steps: 1. 3 Stage Clos, 2. Properties, 3. Select Devices, and 4. Topology Validation. A green notification banner states: "10.20.54.61, 10.20.54.62, 10.20.54.63, 10.20.54.64, 10.20.54.65, 10.20.54.66, 10.20.54.68, 10.20.54.69 added to the fabric successfully".

The main area is titled "Physical Topology" and contains a network diagram. The diagram shows two "Border" devices (10.20.54.68 and 10.20.54.69) at the top, two "Spine" devices (10.20.54.62 and 10.20.54.61) in the middle, and six "Leaf" devices (10.20.54.63, 10.20.54.64, 10.20.54.65, 10.20.54.66, 10.20.54.67, 10.20.54.68) at the bottom. Lines represent connections between these devices.

On the right side, there are two panels:

- Unassigned Devices:** A table with columns: Name, Model, IP, Firmware, LastRef. It lists two devices:




Name	Model	IP	Firmware	LastRef
LS1	3009	10.20.49.118	20.4.3a	26 minutes
LS2	3009	10.20.49.119	20.4.3a	26 minutes
- Assigned Devices (8):** A table with columns: Name, Model, Type, IP, Ports, Firmware. It lists four leaf devices:

Name	Model	Type	IP	Ports	Firmware
Leaf1	3012		10.20.54.63	0/26, 0/28	20.4.3dxcos20.4.3b_230320_05
Leaf2	3012		10.20.54.64	0/26, 0/28	20.4.2
Leaf3	3012		10.20.54.65	0/26	20.4.2

- The discovered topology is displayed. You can select and edit device and fabric configurations directly from the **Physical Topology** or **View Devices** panel as required.
- From the **Device Actions** menu, select **Network Essentials** to modify network essential configurations of the required device ports. For more information, see [Network Essentials](#) on page 129.

All fields marked with an asterisk (*) are required.

Device Actions

-  Delete
-  Firmware Upgrade
-  Network Essentials

Fabric Device Attributes
Links

Hostname (Optional)

Device IP (Optional)

Role (Optional)

Local ASN (Optional)

Loopback ID (Optional)

Cancel
Update

12. Select **Discover Topology** to validate the topology.

The discovered topology is displayed.

The ports or links that are down are marked in red in the topology. To turn a port or link up:

- a. Select the link that is down.
- b. In the **Ports Up** dialog box, configure the ports.
- c. Select **Apply**.
- d. Select **Discover Topology** to validate the topology.

Select **Topology View** () and **List view** () to switch the view between topology and list.

13. To enlarge the topology to the width of the interface, select **Expand** ()

14. To zoom in and out on the topology screen, use the **Zoom** ( ) icons.

15. To scroll through the topology screen, use the **Scroll** () icon.

16. Select **Finish** to configure the topology.

The 3 stage Clos fabric topology is configured.

What to Do Next

Select **View Fabric** or **Proceed to Dashboard** to return to the **Fabrics** page.

Related Topics

[View Fabric Topology](#) on page 124

[Create a Non-Clos Fabric](#) on page 109

[Create a 5 Stage Fabric](#) on page 118

Create a 5 Stage Fabric

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the Fabrics page, select **Create Fabric**.
3. In the **Fabric Name** field, enter a name for the fabric.
4. In the **Fabric Description** field, enter the description for the fabric.

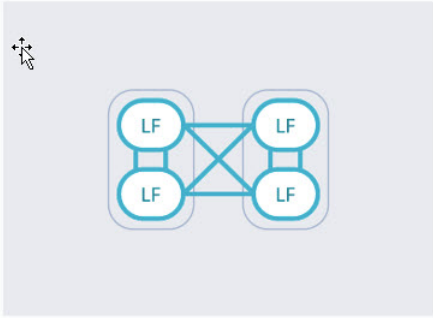
5. Select the **5 Stage Clos** topology.

1 2 3 4 5
5 Stage Clos Properties Select Pods Select Devices Topology Validation

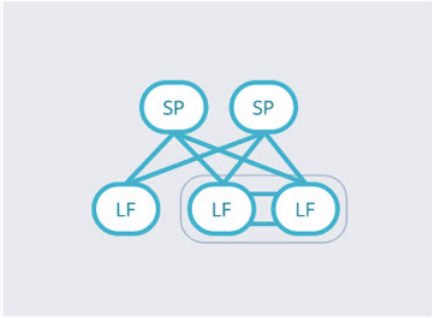
Add Fabric Name and Select Type

Fabric Name *

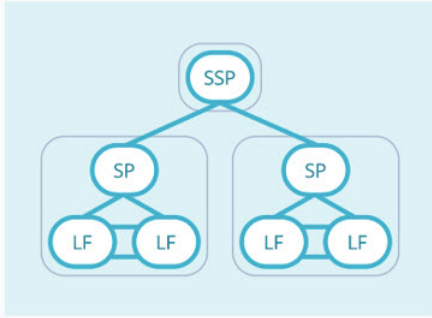
Fabric Description (Optional)



Non Clos
Non CLOS topology involves n (1 to 4) number of racks interconnected to each other. Rack consists of a pair of leaf or a pair of border leaf



3 Stage Clos
3-Stage CLOS topology involves Spine Layer and Leaf/Border Leaf Layer. Leaf/Border Leaf can be single-homed or dual-homed.



5 Stage Clos
5-Stage CLOS topology involves Super Spine Layer, Spine Layer and Leaf/Border Leaf Layer. Leaf/Border Leaf can be single-homed or dual-homed.

Next

6. Select **Next**.

7. In the **Fabric Properties** page, complete the fields as required or select the **Auto Config Mode** check box to use the default fabric settings.

1 2 3 4 5
5 Stage Clos Properties Select Pods Select Devices Topology Validation

Fabric Properties

Auto Config Mode ⓘ

✕ OVERLAY

All fields marked with an asterisk (*) are required.

Enable Overlay

VNI Auto Map

Broadcast Local Bias

✕ IP RANGE

All fields marked with an asterisk (*) are required.

Fabric Link IP Range *

Loopback IP Range *

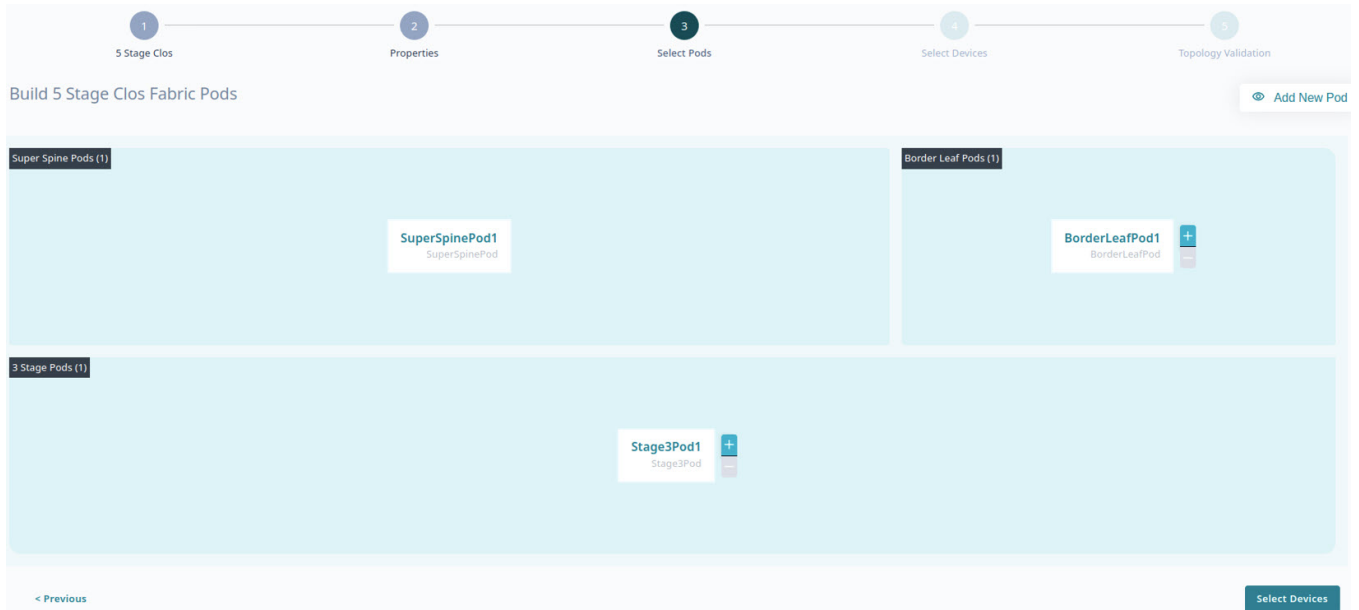
MCT Link IP Range *

✕ BFD (Bidirectional Flow Detection)

< Previous Create Fabric

8. Select **Create Fabric**.

9. In the **Build 5 Stage Clos Fabric Pods** page, select **+** or **-** to add new 3 stage or border leaf pods.

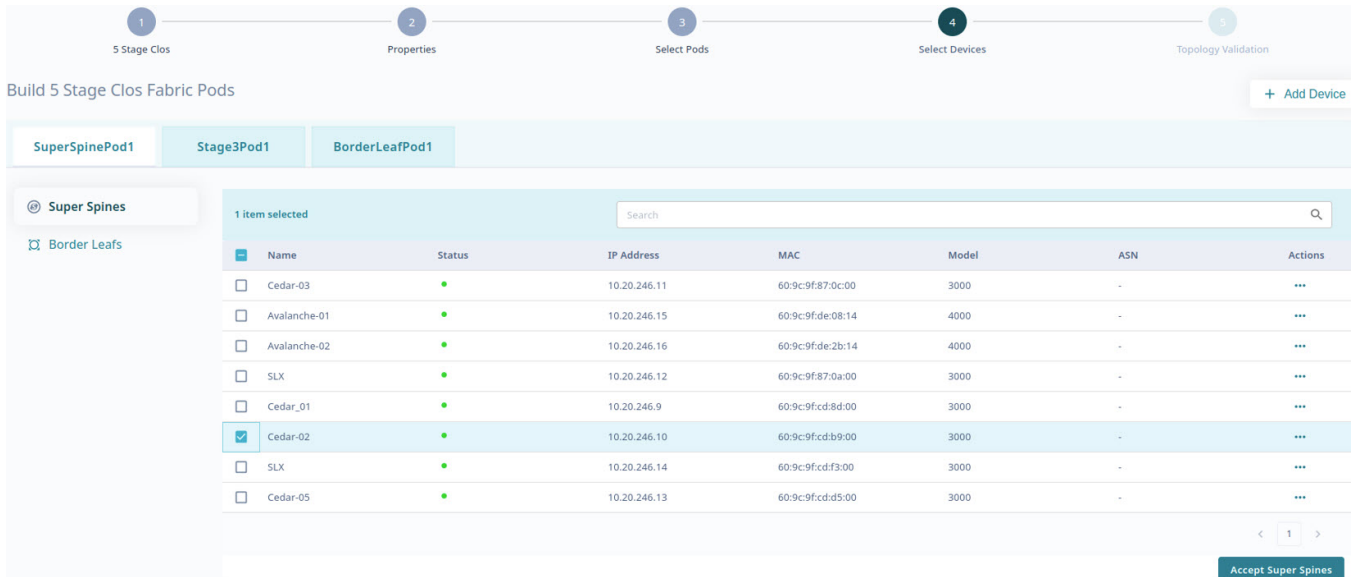


- Alternatively, you can select **Add New Pod**.
- One pod of each type is available in the UI by default and the pod names are auto-generated. For example: SuperSpinePod1, Stage3Pod1, and BorderLeafPod1.

10. Select **Select Devices** to add new devices.

11. Select the check boxes of the required devices from the following tabs:

- **SuperSpinePod1**
- **Stage3Pod1**
- **BorderLeafPod1**



12. Select **Accept Super Spine Pods**, **Accept Spines**, or **Accept Border Leaf Pods** as applicable.

13. Select **Accept All the Pods** to add all devices to the topology.

The screenshot shows the 'Physical Topology' discovery process in the ExtremeCloud Orchestrator GUI. The interface is divided into several sections:

- Progress Bar:** Shows five stages: 1. 5 Stage Clos, 2. Properties, 3. Select Pods, 4. Select Devices, and 5. Topology Validation. Stage 5 is currently active.
- Notification:** A green banner at the top left states 'Topology discovered successfully'.
- Physical Topology:** The main area displays a network diagram with nodes for BorderLeafPod1, SuperSpinePod1, and StageSpinePod1. A search bar for IP addresses and an 'Add New Pod' button are visible.
- Unassigned Devices:** A table on the right lists devices that have not been assigned to the topology. The table has columns for Name, Model, IP, Firmware, and LastRef.
- Assigned Devices (7):** A table on the right lists devices that have been assigned to the topology. The table has columns for Name, Model, Type, IP, Ports, and Firmware.

- The discovered topology is displayed. You can select and edit the device configuration directly from the **Physical Topology** or **View Devices** panel as required.
- You can select devices in rack and update fabric specific configurations such as ASN, VTEP Loopback ID, and Loopback ID. The attributes vary by device role.
- From the **Device Actions** menu, select **Network Essentials** to modify network essential configurations of the required device ports. For more information, see [Network Essentials](#) on page 129.

All fields marked with an asterisk (*) are required.

Device Actions

- Delete
- Firmware Upgrade
- Network Essentials

Fabric Device Attributes **Links**

Hostname (Optional)

Device IP (Optional)

Role (Optional)

Local ASN (Optional)

Loopback ID (Optional)

Select **Topology View** () and **List view** () to switch the view between topology and list.

14. To enlarge the topology to the width of the interface, select **Expand** ().

15. To zoom in and out on the topology screen, use the **Zoom** ( ) icons.

16. To scroll through the topology screen, use the **Scroll** () icon.

17. Select **Finish** to configure the topology.

The 5 stage Clos fabric topology is configured.

What to Do Next

Select **View Fabric** or **Proceed to Dashboard** to return to the Fabrics page.

Related Topics

[View Fabric Topology](#) on page 124

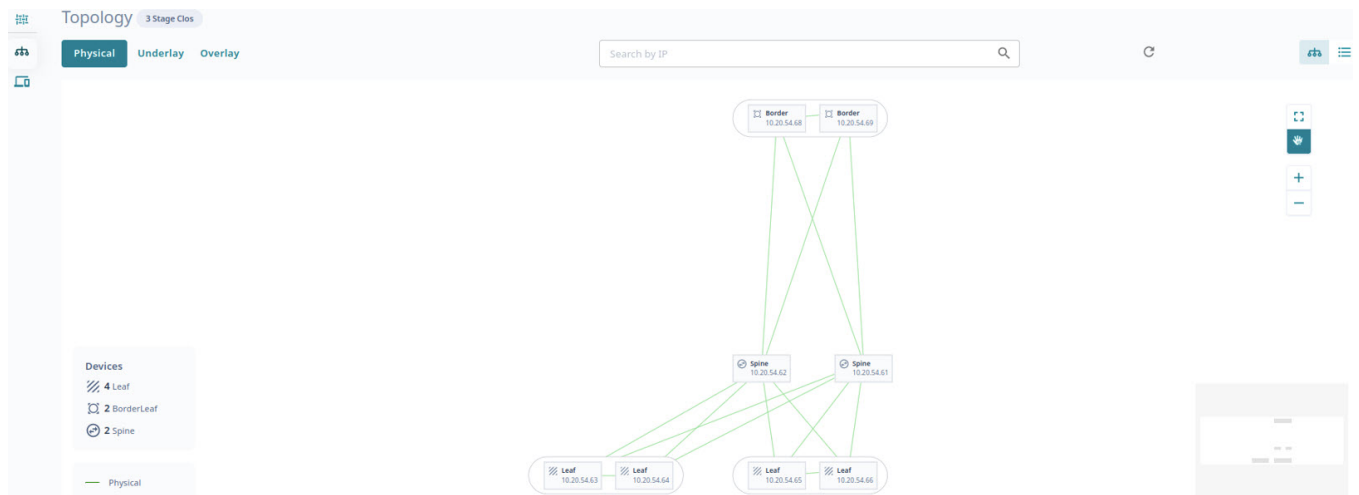
[Create a Non-Clos Fabric](#) on page 109


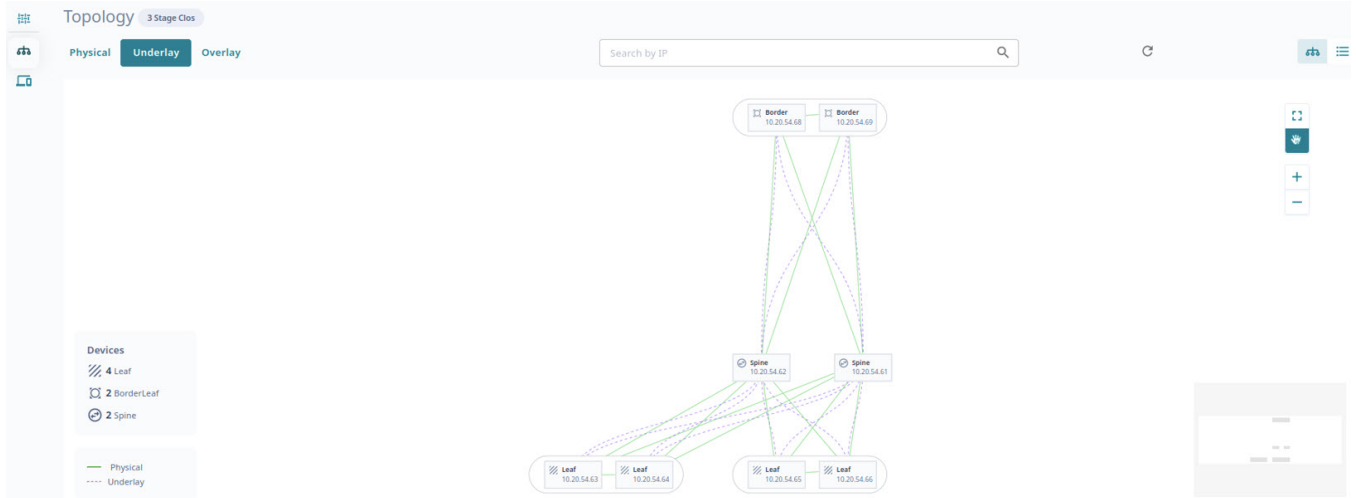
[Create a 3 Stage Clos Fabric](#) on page 113

View Fabric Topology

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the **Fabrics** page, click anywhere in the required fabric row except the Actions column (•••) to proceed to the fabric Topology page.
3. Select the required topology tab.
 - **Physical**: Represents physical connections of the fabric devices
 - **Underlay**: Represents BGP sessions between the fabric devices
 - **Overlay**: Represents the overlay (VXLAN) tunnel state between leaf or border-leaf devices





Select **Topology View** (🔗) and **List view** (☰) to switch the view between topology and list.

The screenshot shows the 'List view' of the network topology. It displays a table with columns for Encap Type, Tunnel Type, Source Leaf IP, Destination Leaf IP, Source VTEP IP, Destination VTEP IP, Admin State, and OPER State. The table shows 6 results for the topology.

Encap Type	Tunnel Type	Source Leaf IP	Destination Leaf IP	Source VTEP IP	Destination VTEP IP	Admin State	OPER State
▼ 10.20.54.64,10.20.54.63							
vxlan	unicast	10.20.54.64,10.20.54.63	10.20.54.65,10.20.54.66	172.31.254.146	172.31.254.97	up	up
vxlan	unicast	10.20.54.64,10.20.54.63	10.20.54.68,10.20.54.69	172.31.254.146	172.31.254.3	up	up
▼ 10.20.54.65,10.20.54.66							
vxlan	unicast	10.20.54.65,10.20.54.66	10.20.54.64,10.20.54.63	172.31.254.97	172.31.254.146	up	up
vxlan	unicast	10.20.54.65,10.20.54.66	10.20.54.68,10.20.54.69	172.31.254.97	172.31.254.3	up	up

Download Fabric Inventory


Procedure


1. In the Navigation menu, select **Fabrics**.

2. Select  **Download**.
A file in .csv format is downloaded to your device.

Delete Fabric

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the **Fabrics** page, select **Delete** () from the Actions column (**⋮**) for the fabric you want to delete.



3. Select **Confirm** when prompted.

Download Health Report

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the **Fabrics** page, select **Download Health Report** from the Actions column (**⋮**) for the required fabric.
The fabric health report is downloaded to your device.

Fabric Actions

You can perform the following tasks from the **Fabric Actions** menu on the **Fabrics** page.

- Edit
- Delete
- Firmware Upgrade
- Clone

Edit Fabric Topology

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the **Fabrics** page, click anywhere in the required fabric row except the Actions column (•••) to proceed to the fabric Topology page.
3. In the **Fabric Actions** menu, select **Edit**.
The fabric topology is displayed.
4. In the **Topology** page, do the following as required:
 - Select **Settings** to edit appropriate fabric properties.

- Select **Devices** to edit devices in the topology.

The screenshot shows the 'Editing Stage3_CLOS' topology in the ExtremeCloud Orchestrator GUI. The main window displays a network diagram with 'Border' and 'Leaf' devices connected. A 'Devices' panel on the right lists unassigned and assigned devices with columns for Name, Model, IP, Firmware, and LastRef. The 'Assigned Devices' section shows 6 devices with their respective details.

Name	Model	IP	Firmware	LastRef
Cedar_01	3001	10.20.246.9	18s.1.03	58 minutes
Freedom-03	3001	10.20.246.19	18s.1.03	19 minutes
Freedom-05	3001	10.20.246.21	18s.1.01a	58 minutes
Freedom_06	3001	10.20.246.22	18s.1.01a	19 minutes
Freedom-04	3001	10.20.246.20	18s.1.03	18 minutes

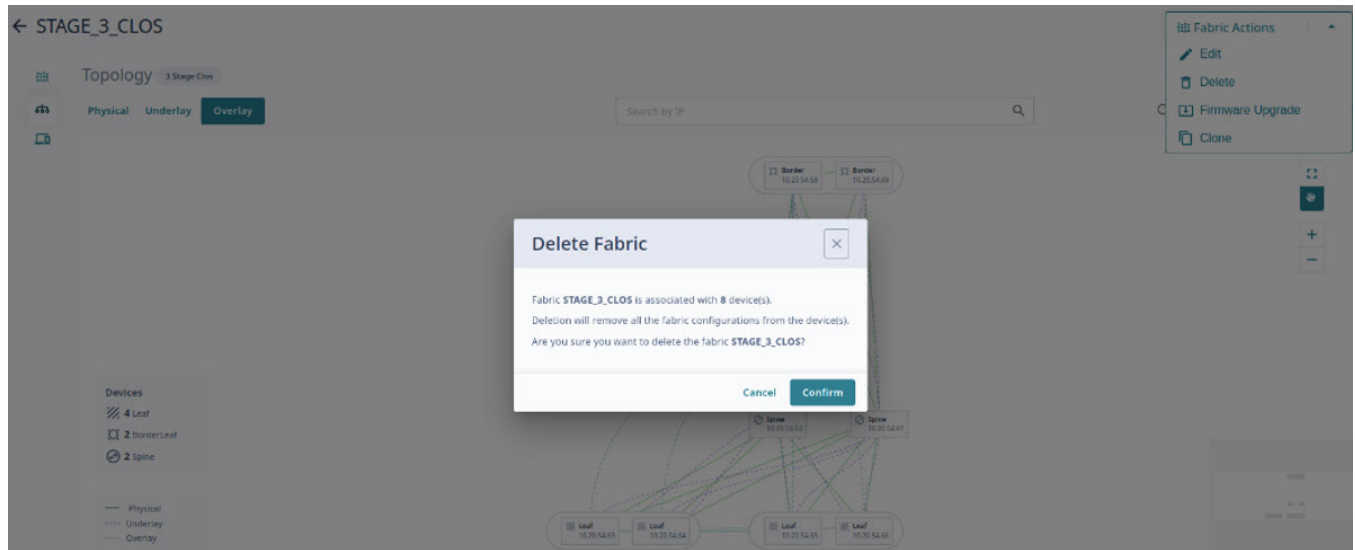
- Alternatively, you can select a device directly from the topology to access **Device Information** and edit **Fabric Device Attributes** as required.
 - In the **Device Actions** menu, select **Network Essentials** to modify network essential configurations of the required device ports. For more information, see [Network Essentials](#) on page 129.
5. Select **Discover Topology** to verify the links in the topology.
 6. Select **Update Fabric** to update the fabric.

Delete a Fabric from the Fabric Actions Menu

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the **Fabrics** page, click anywhere in the required fabric row except the Actions column (•••) to proceed to the fabric Topology page.

3. Select **Delete** from the **Fabric Actions** menu.



- If devices are present in the fabric, a message is displayed to use the **Force** option to delete the fabric.
 - Proceed to the next step to force delete the fabric, else skip the next step.
4. Select the **Force** check box.
 5. Select **Confirm** when prompted.

Network Essentials

XCO 3.2.0 supports the following network essential configurations that are required for creating and configuring fabric networks:

- Admin State (up/down)
- MTU (L2/Ipv4/Ipv6)
- Speed
- Breakout
- FEC (Forward Error Correction)
- Link Error
- RME (Redundant Management Ethernet)

Configure Network Essentials

Procedure

1. In the Navigation menu, select **Fabrics**.
2. In the **Fabrics** page, click anywhere in the required fabric row except the Actions column (•••) to proceed to the fabric Topology page.
3. Select the required device in the topology.
4. In the Fabric Actions menu, select **Network Essentials**.
 - Network essential configuration of all ports in the selected device is displayed.

- Alternatively, you can access **Network Essentials** configurations from the Actions column (☰) in the **Device Inventory** and **Fabrics** pages.

5. Edit the required ports.

Host Name: Device 1
IP Address: 192.168.1.11
Model: SLX 9250

Ports

0/3 X

Search

<input type="checkbox"/>	Name	Admin State	Speed	Breakout	L2MTU	IPv4MTU	IPv6MTU	FEC	Dampening	Damp. Toggle Threshold	Damp. Sampling Time	Damp. Wait Time	RME
<input type="checkbox"/>	0 / 1	Down	100Gbps	None	1500	1500	1500	rs-fec	True	1000	1222	300	True
<input type="checkbox"/>	0 / 2	Down	100Gbps	None	1500	1500	1500	rs-fec	True	1000	1222	300	True
<input checked="" type="checkbox"/>	0 / 3	Down	100Gbps	None	1500	1500	1500	rs-fec	True	1000	1222	300	True
<input type="checkbox"/>	0 / 4	Down	100Gbps	None	1500	1500	1500	rs-fec	True	1000	1222	300	True

Showing 1 - 13 of 46 results

Reset Apply Network Essentials

6. Select **Apply Network Essentials**.

Firmware Upgrade

XCO supports firmware download and upgrade across all devices of the fabric.

Before You Begin

- Register firmware host. For more information, see [Register Firmware Host](#) on page 79.
- XCO determines the grouping of devices for firmware download to achieve least traffic disruption when upgrading a fabric with active traffic.
- You can select single or multiple devices in the fabric for firmware upgrade.

About This Task



Note

As a best practice, do not change the target firmware version file name and the directory name.

Procedure

- In the Navigation menu, select **Fabrics**.

- In the **Fabrics** page, click anywhere in the required fabric row except the Actions column (⋮) to proceed to the fabric Topology page.

Showing 1 - 3 of 3 results

Name	Health	Type	Stage	Status	Leafs	Border Leafs	Spines	Super Spines	Actions
Stage3_CLOS	Healthy	Clos	3 Stage	Configure-success	2	2	2	n/a	⋮
Stage_5_CLOS	Critical	Clos	5 Stage	Configure-success	2	2	2	1	⋮
default	Healthy	Clos	3 Stage	Created	0	0	0	n/a	⋮

- (Optional) In the **Topology** page, go to **Device Management** (🖨️) to select the check boxes of the required devices for firmware upgrade. Skip this step to upgrade all devices in the fabric.

- Select **Firmware Upgrade** from the Fabric Actions menu.

Upgrade Device Firmware ✕

All fields marked with an asterisk (*****) are required.

Host *

10.31.2.101

Absolute Path *

/buildsjc/sre/fusion/Nightly/raphael/slxos20.4.3a/LATES

Provide absolute path where the firmware bundle is stored

Firmware Upgrade Options

Download
Download the firmware and prepare the device for firmware upgrade.

Activate
Activate firmware for which the device is already prepared for.

Download and Activate
Download the firmware, prepare the device and activate the firmware upgrade.

Auto Commit ⓘ

Selected Devices (6)

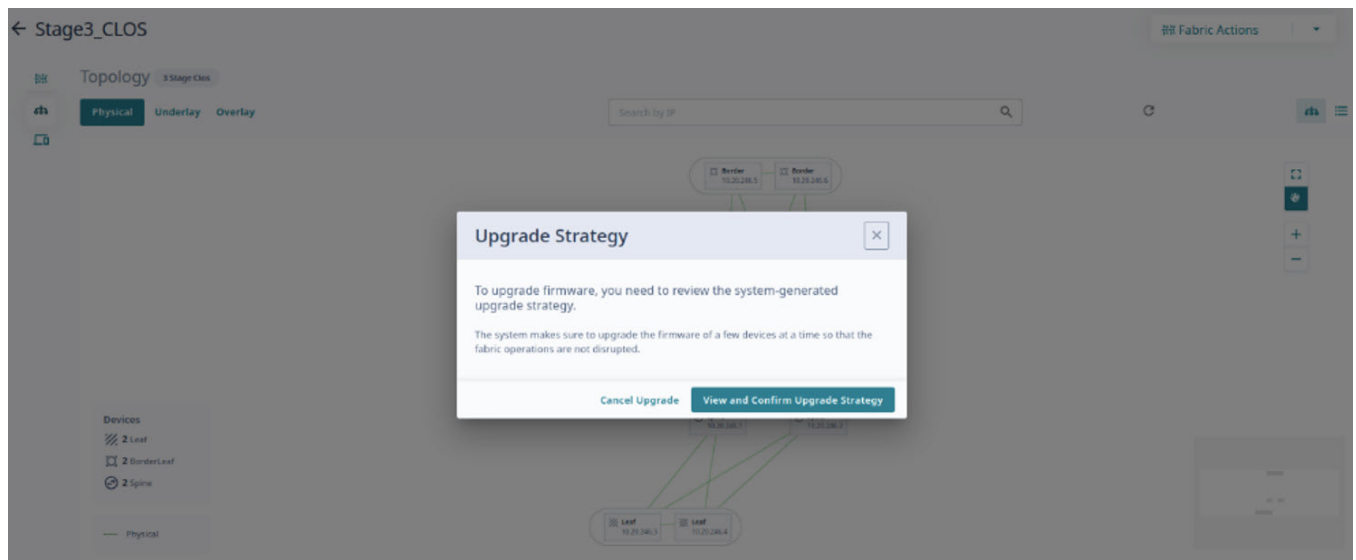
Showing 1 - 6 of 6 results

Model	Role	Firmware	IP Address
3012	Spine	20.4.3slxos20	10.20.246.1
3009	BorderLeaf	20.4.3slxos20	10.20.246.5
3009	BorderLeaf	20.4.3slxos20	10.20.246.6
3012	Spine	20.4.3slxos20	10.20.246.2
3012	Leaf	20.4.3slxos20	10.20.246.4
3012	Leaf	20.4.3slxos20	10.20.246.3

Cancel
Upgrade Firmware

- In the **Host** field, provide the IPv4 or IPv6 address of the firmware host server.
- In the **Absolute Path** field, provide the firmware file path.
- Select **Download and Activate**.
- (Optional) Enable or disable **Auto Commit** as required.

9. If Auto Commit is disabled, do the following as required on the fabric devices pending for commit:
 - In the Fabric Actions menu, select **Commit Firmware Upgrade** to commit the firmware.
 - In the Fabric Actions menu, select **Restore Firmware Upgrade** to restore the firmware.
10. Select **Upgrade Firmware**.
 - The list of devices in the fabric is displayed.
 - The LLDP links of the devices in the fabric might go down during firmware download as devices reload and will be in the maintenance mode. This is reflected in the fabric topology view as "No physical links discovered".
 - You are prompted to review the system generated upgrade strategy to minimize traffic disruption to the active fabric.
11. Select **View and Confirm Upgrade Strategy** to review and approve the device upgrade sequence.



Fabric Devices Upgrade Strategy
✕

Device Upgrade Sequence

- Upgrade all leaf devices on the Left
- Upgrade all leaf devices on the Right
- Upgrade half of all Border Leaf devices
- Upgrade other half of all Border Leaf devices
- Upgrade Spine 1
- Upgrade Spine 2
- Commit Upgrade manually

⚠ As you have selected "Auto Commit" to be "OFF", the firmware will not be committed on secondary partition on the device. The device will be in "Commit Pending" state. Once you are satisfied with the upgraded firmware, commit the upgrade by using "Commit Upgrade" from the action menu.

⏪ ⏩

⊙ ⊙
⊙ ⊙ ⊙ ⊙ ⊙ ⊙
⊙ ⊙ ⊙ ⊙

Auto Commit ⚠ To finalize your Upgrade process please Commit Upgrade later, or turn on Auto commit now
Cancel Confirm Upgrade

12. Select **Confirm Upgrade**.

- The list of devices in the fabric along with upgrade status is displayed.
- The device upgrade status indicates various stages such as download, active, and commit. The user interface also provides updates such as the number of devices undergoing upgrade, waiting for upgrade, activation and commit pending, commit upgrade, restore upgrade, and upgrade success.

← Stage3_CLOS
Fabric Actions

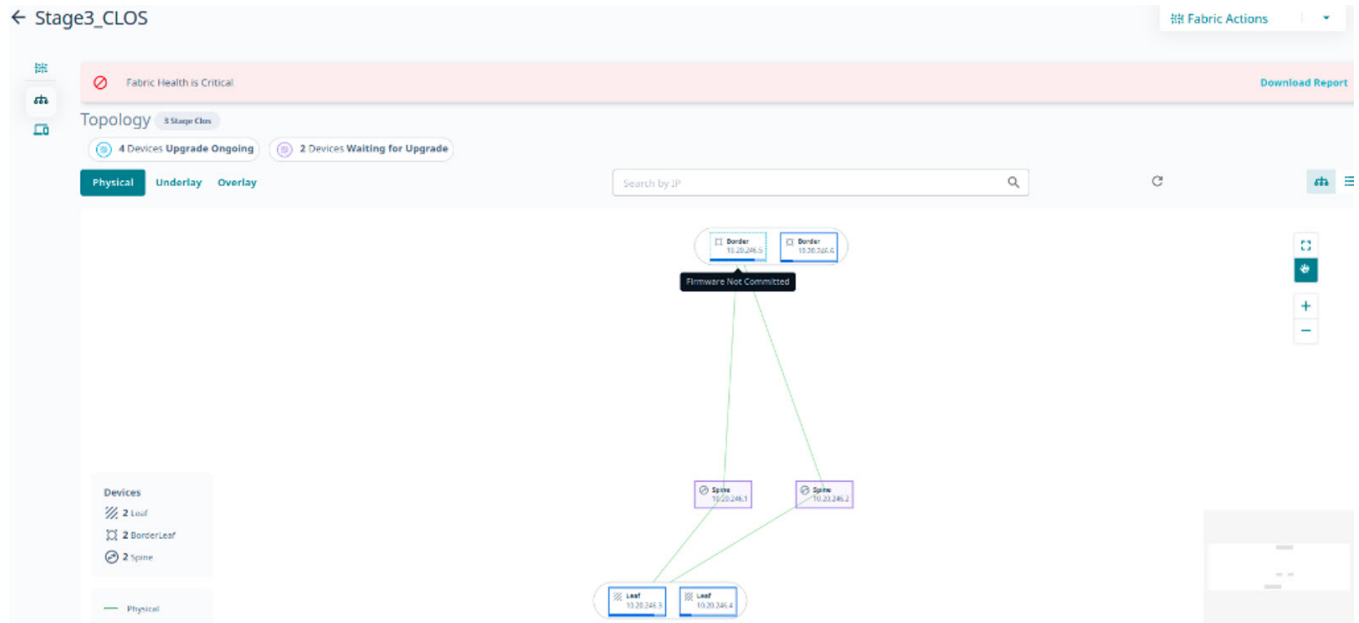
Fabric Health is Critical
Download Report

Devices

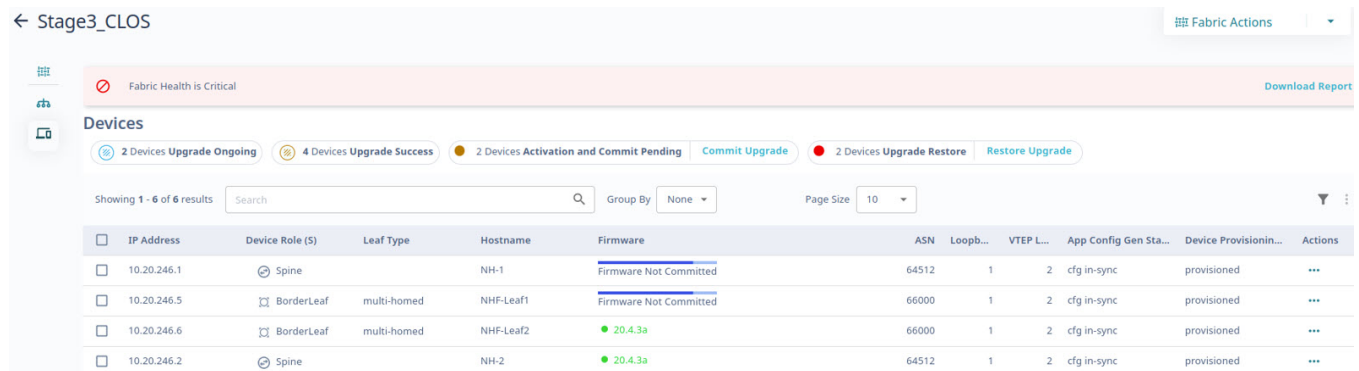
2 Devices Upgrade Ongoing
4 Devices Waiting for Upgrade

Showing 1 - 6 of 6 results Group By: None Page Size: 10

IP Address	Device Role (5)	Leaf Type	Hostname	Firmware	ASN	Loopb...	VTEP L...	App Config Gen Sta...	Device Provisionin...	Actions
<input type="checkbox"/> 10.20.246.1	⊙ Spine		NH-1	20.4.3s1xos20.4.3a_230218_1918	64512	1	2	cfg in-sync	provisioned	⋮
<input type="checkbox"/> 10.20.246.5	⊙ BorderLeaf	multi-homed	NHF-Leaf1	Device Reload Started	66000	1	2	cfg in-sync	provisioned	⋮
<input type="checkbox"/> 10.20.246.6	⊙ BorderLeaf	multi-homed	NHF-Leaf2	20.4.3s1xos20.4.3a_230218_1918	66000	1	2	cfg in-sync	provisioned	⋮
<input type="checkbox"/> 10.20.246.2	⊙ Spine		NH-2	20.4.3s1xos20.4.3a_230218_1918	64512	1	2	cfg in-sync	provisioned	⋮
<input type="checkbox"/> 10.20.246.4	⊙ Leaf	multi-homed	NH-Leaf2	20.4.3s1xos20.4.3_221117_0600	65000	1	2	cfg in-sync	provisioned	⋮
<input type="checkbox"/> 10.20.246.3	⊙ Leaf	multi-homed	NH-Leaf1	Firmware Download And Maintenance Mode Ena	65000	1	2	cfg in-sync	provisioned	⋮



13. Select **Commit Upgrade** to commit pending devices.
Alternatively, select **Commit Firmware Upgrade** from the Fabric Actions menu.
14. Select **Commit Upgrade** when prompted.
The devices are upgraded to the downloaded firmware version.



Related Topics

- [Register Firmware Host](#) on page 79
- [View Registered Firmware Hosts](#) on page 80
- [Change a Firmware Host](#) on page 81
- [Delete a Firmware Host](#) on page 81

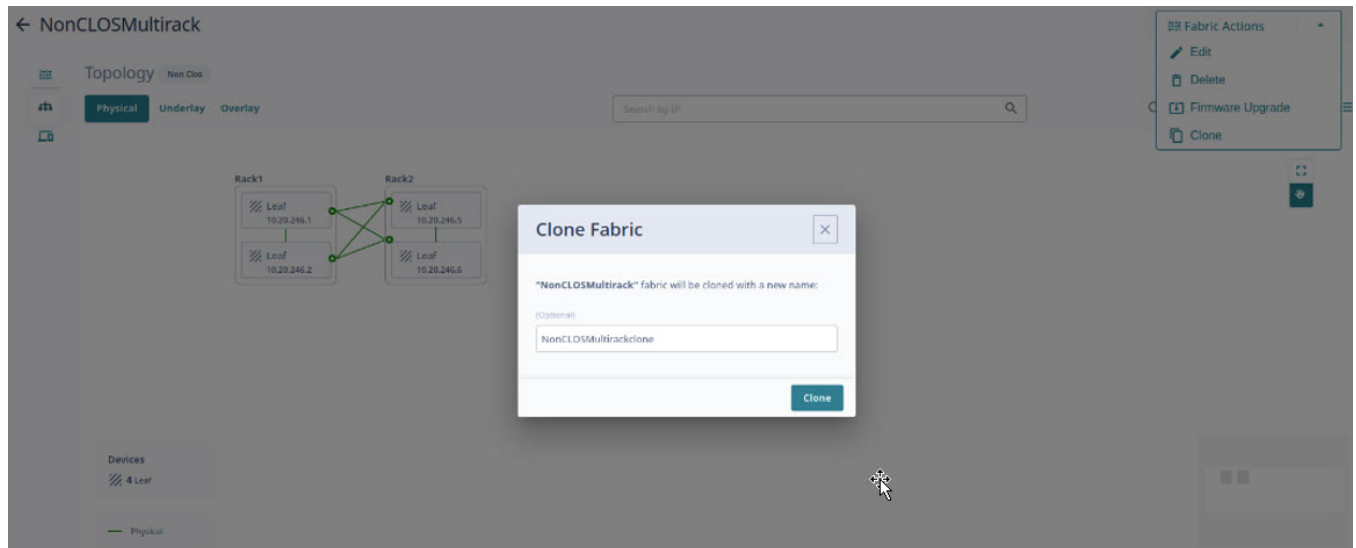
Clone a Fabric

You can clone (copy) a fabric to create a new fabric with the same or similar topology.

Procedure

1. In the Navigation menu, select **Fabrics**.

2. In the **Fabrics** page, click anywhere in the required fabric row except the Actions column (•••) to proceed to the fabric Topology page.
3. In the Fabric Actions menu, select **Clone** (📄).



4. Enter a name for the new fabric.
5. Select **Clone**.



FAQs

Where are Visibility skill logs located?

Debug logs: `/var/log/efa/evm/evm-server.log`

Error logs (with panic trace): `/var/log/efa/evm/evm-server_err.log`

Where are the Inventory Service logs located?

Debug logs: `/var/log/efa/inventory/inventory-server.log`

Error logs (with panic trace): `/var/log/efa/inventory/inventory-server_err.log`

Where are the Installer logs located?

`/var/log/efa/installer/<installer_.....log>`

All installation failures are reported in this log.

What are some common reasons for installation failures?

- The operating system version is incorrect.
- The amount of available hard disk space is insufficient. At least 50 GB should be available.
- In a multi-node installation, the operating system and clock do not match. Or, both nodes have the same host name.

Why does the web user interface not load on the browser?

The most probable reason is that TCP port 443 is blocked through a firewall. Unblocking this port should enable the UI to be loaded.

What are some common reasons for XCO log-in failures?

- The user credentials are entered incorrectly.
- TACACS+ or LDAP is not reachable or not configured correctly for the `xco-role`.
- The `xco-role` in TACACS+ or LDAP is not mapped to a predefined role such as NetworkOperator or SystemAdmin.

Where are authentication failures captured?

Debug logs: `/var/log/efa/auth/auth-server.log`

Error logs (with panic trace): `/var/log/efa/auth/auth-server_err.log`

What are possible reasons for device registration failures?

- The device is not reachable from ExtremeCloud Orchestrator.
- The device credentials are incorrect.
- The HTTPS, SSH, NETCONF, or GNMI ports are blocked.
- The device versions are not supported.
- The device has exceeded the maximum limit on SSH connections. Free up some existing connections that are used by other tools and try to register again.

Why is there a delay in loading the dashboard or statistics in the web UI?

It can take from 20 seconds to 1 minute to load live statistics from a device.

Why is the device configuration blocked from the web UI?

The device can have missed a heartbeat and subsequently transitioned to a degraded state. The device should be accessible when connectivity is restored.

What are possible reasons for configuration failures?

- The XCO user does not have permission to make changes to the device.
- The web UI reports validation errors or errors received from the device.
- The credentials used for device registration do not have permission to make changes to the device.

How do I check that all services are up and running?

Run the following command on the XCO device:

```
k3s kubectl get pods -n efa
```

The following is sample output.

NAME	READY	STATUS	RESTARTS	AGE
efa-api-docs-84cwl	1/1	Running	0	20d
ui-service-54dbbb47fd-vzfrw	1/1	Running	0	20d
gosystem-service-dw4vj	1/1	Running	0	20d
rabbitmq-4tgsv	1/1	Running	0	20d
gorbac-service-j6lp8	1/1	Running	0	20d
goevm-service-sjckq	1/1	Running	0	20d
gonotification-service-grvx9	1/1	Running	0	20d
gofaultmanager-service-nstwf	1/1	Running	0	20d
goauth-service-qffvs	1/1	Running	0	20d
goraslog-service-s7mwt	1/1	Running	0	20d
goinventory-service-q5sdl	1/1	Running	0	20d

Why are the device syslogs not visible?

Other tools that are registered with the device could have exceeded the maximum limit for syslogs. Free up any stale syslog entries on the device and then re-register the device.

How to collect the SupportSave data for troubleshooting?

Run the following command on the XCO device:

```
efa system supportsave
```

The following is a sample output.

```
SupportSave File Location: /var/log/efa/efa_2022-11-17T18-40-41.008.logs.zip  
--- Time Elapsed: 21.584259642s ---
```

To collect the SupportSave data using the XCO GUI, see [Support Save](#) on page 25.