ExtremeCloud™ Orchestrator v4.0.2 Release Notes

Resolved and Open Issues

9039317-01 Rev AA
April 2026

Legal Notice

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

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

Trademarks

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

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

For additional information on Extreme Networks trademarks, see: https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks

Open Source Declarations

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

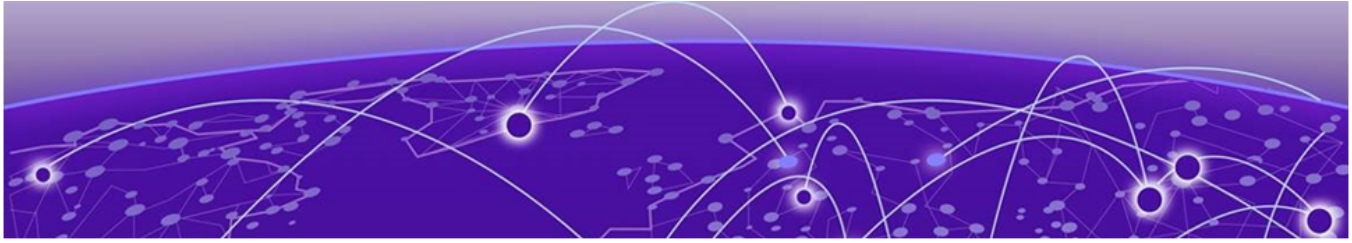

Abstract

ExtremeCloud™ Orchestrator v4.0.2 Release Notes supports lifecycle management capabilities, supporting advanced fabric provisioning, multi-tenant networking, and automation features. This release resolves issues through defect fixes.

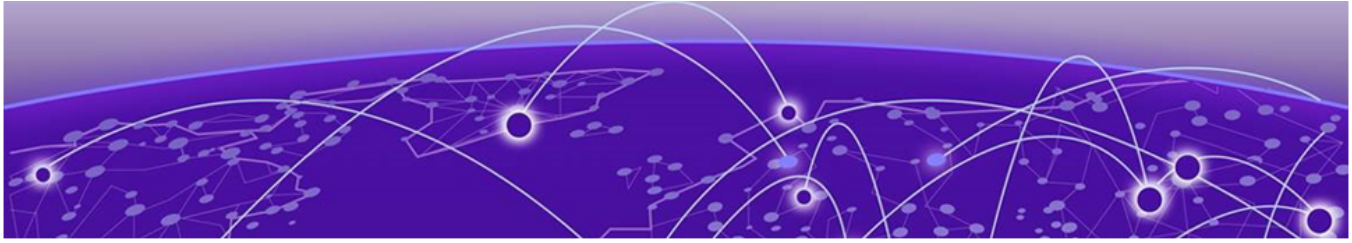

Release Notes

New In This Release on page 5
XCO Extreme OS ONE Integration on page 6
Supported Platforms and Deployment Models for Fabric Skill on page 14
XCO Upgrade Prerequisites on page 17
Open Issues on page 18

New In This Release

ExtremeCloud Orchestrator 4.0.2 introduces the following features and enhancements, and resolves issues through defect fixes.

Note

- In release 3.2.0 and later, Extreme Fabric Automation (EFA) is referred to as ExtremeCloud Orchestrator (XCO). The terms EFA and XCO refer to the same product and are used interchangeably.
- In release 4.0.2, the "ID" field in the Request Body for the RegisterTrapSubscriber REST API is read-only. When registering a Trap Subscriber using REST API, the "ID" field must not be included in the request body. This change does not affect Trap Subscriber registration using CLI.

Table 1: Features and Enhancements

Feature	Description
XCO Auditd Enhancement	Supports Auditd (Linux Auditing System) on the Host OS to record and track system-level security events. This feature provides detailed logging of security-relevant activities, enabling administrators to monitor potential security incidents and maintain policy compliance. For more information, see *ExtremeCloud Orchestrator v4.0.2 CLI Administration Guide*.
BGP Max-Path Support	Adds BGP maximum-paths support in the SLX fabric, enabling Equal-Cost Multi-Path (ECMP) routing for BGP routes. This feature allows traffic to be load-balanced across multiple equal-cost paths throughout the fabric. For more information, see *ExtremeCloud Orchestrator v4.0.2 CLI Administration Guide*. For more information, see *ExtremeCloud Orchestrator v4.0.2 GUI Administration Guide*.
Link Add (LA) and Link Delete (LD) Enhancement	Ensures that when LA to LD or LD to LA events occur in sequence, devices transition to cfg-in-sync state immediately, without waiting for DRC or Fabric Configure. For more information, see *ExtremeCloud Orchestrator v4.0.2 CLI Administration Guide*.
Notification Subscriber With FQDN	Configure a notification subscriber using a Qualified Domain Name (FQDN). For more information, see *ExtremeCloud Orchestrator v4.0.2 CLI Administration Guide*.
RELP (Syslog) and Webhook Remote Server Secure Notifications Configuration	Supports secure configuration of remote syslog and webhook. For more information, see *ExtremeCloud Orchestrator v4.0.2 CLI Administration Guide*.

Table 1: Features and Enhancements (continued)

Feature	Description
L3 Extension Configuration in Single Rack Non-Clos Deployment	Controls single rack fabric behavior on SLX. For more information, see *ExtremeCloud Orchestrator v4.0.2 CLI Administration Guide*.
Firmware Upgrade	Updated prerequisites for upgrading to XCO 4.0.2 without Ubuntu 20.04 support. For more information, see *ExtremeCloud Orchestrator v4.0.2 Deployment Guide*.
New/Modified Commands	New commands: • efa policy debug drift-reconcile Modified commands: • efa fabric setting • efa inventory device interface set-fec • efa inventory device interface set-link-error-disable • efa system feature update • efa tenant epg create • efa tenant epg update • efa tenant vrf update For more information, see *ExtremeCloud Orchestrator v4.0.2 Command Reference*.

XCO Extreme OS ONE Integration

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

Note

- The features related to XCO OS ONE and Extreme 8730 support in XCO 4.0.0 are released as Control Released Features.
- For comprehensive feature descriptions, CLI command references, administrative procedures, and API documentation, contact Extreme Networks.

The following features are included as Control Released Features in XCO 4.0.0:

- XCO OS ONE Integration: Lifecycle management support for XCO OS ONE-based devices.
- Extreme 8730 Support: Inventory and fabric integration for the new platform.

- Secure Settings via GNMI: Centralized enforcement of SSH, TLS, and password policies.
- CLI and API Enhancements: Extended command coverage for device and fabric operations.

Note
Config Drift Tracking is not supported.

Table 2: Supported OS ONE features

Feature	Description
Fabric	
Non-Clos Fabric Support	XCO automates both Clos and non-Clos environments. The following are supported for non-Clos fabrics: • Automated provisioning of Layer 2 and Layer 3 services • Tenant-aware networking, including VLANs and VRFs • Simplified configuration workflows via CLI and REST APIs • MCT (Multi-Chassis Trunk) for redundancy XCO uses microservices running on a lightweight Kubernetes cluster (K3s) to orchestrate these functions, whether deployed on TPVM (Third-Party Virtual Machine) or external VMs.
Tenant	
Tenant PO	A port channel, also known as a Link Aggregation Group (LAG) is a communication link between devices. For creation of single-homed or multi-homed port channel for OS ONE devices, you can specify the following parameters: LACP negotiation, LACP timeout, Port Port channel number, MTU, Number of active links required. Speed is not required for port channel creation unlike SLX devices.
BD based L3 EPG Support (Auto-VNI) & VLAN based L3EPG Auto-VNI	An endpoint group is a logical group of endpoints or devices connected to the network. These endpoints are used to deploy Layer-2 or Layer-3 services over the fabric. To create Layer-3 service, you can specify parameters: Group name IP address, port channels, switchport mode, CTAG range, associated VRF, bridge domain, and neighbour discovery preferences.

Table 2: Supported OS ONE features (continued)

Feature	Description
Auto VNI Map Fabric with BD-Based Tenant and L2 EPG	Enables automatic VXLAN Network Identifier (VNI) mapping for tenants in a non-Clos fabric, using Bridge Domain (BD)-based segmentation and Layer 2 Endpoint Groups (EPGs). The following features are supported: • Auto VNI Mapping: ◦ For VLAN-based tenants ◦ For BD-based tenants • BD-Based Tenants: ◦ Multiple tenants can share the same VLAN ◦ Each VLAN is mapped to a unique BD, which then maps to a unique VNI. • L2 EPG: ◦ Allows grouping of endpoints at Layer 2 ◦ Facilitates policy enforcement and traffic segmentation within the same subnet. • Multi-Tenant: ◦ Tenants can have overlapping VLANs (ctags) ◦ Each tenant gets a unique L2VNI and isolated network space. Ideal for multi-tenant data centers where isolation, scalability, and automation are critical.
Auto VNI Map Fabric - VLAN Based Tenant, L2 EPG	Enables automatic VNI assignments for VLAN-based tenants in a Clos fabric, with support for Layer 2 Endpoint Groups (EPGs) as part of the Tenant Service integration in XCO. The following features are supported: • Auto VNI Mapping: ◦ For VLAN-based tenants, the VNI is statically derived from the VLAN ID ◦ Simplifies VXLAN configuration by eliminating manual VNI assignment. • L2 EPG: ◦ Groups endpoints within the same VLAN/subnet ◦ Enables policy enforcement and traffic segmentation at Layer 2 • Multi-Tenant Isolation: ◦ Each tenant gets a unique L2VNI ◦ Prevents VLAN overlap across tenants by enforcing one-to-one VLAN-to-VNI mapping ◦ Ideal for multi-tenant data centers where VLAN-based segmentation is preferred, and automation of VXLAN overlays is needed for scalability and consistency.

Table 2: Supported OS ONE features (continued)

Feature	Description
Virtual Routing and Forwarding (VRF)	Enables Tenant Service to integrate with VRF instances, allowing for multi-tenant Layer 3 segmentation within a non-Clos fabric. The following features are supported: • Tenant-Specific VRFs: ◦ Allows assignment of one or more VRFs to each tenant for isolated routing domains. ◦ Prevents route leakage and ensures traffic separation • Distributed and Centralized Routing: ◦ Allows configuration of VRFs for distributed routing on leaf switches or centralized routing on border nodes. • Lifecycle Management: ◦ Create, update, show, and delete VRFs via CLI or REST APIs ◦ Integrated with XCO inventory and policy engine ◦ Ideal for multi-tenant data centers or service provider environments where Layer 3 isolation is required between tenants. ◦ Scalable and secure routing without needing multiple physical routers.
Inventory	
Network Essentials: NTP	Ensures all devices in the fabric maintain synchronized time—a critical requirement for logging, monitoring, and security. The following features are supported: • Inventory Integration: XCO tracks and manages NTP server configurations across devices in its inventory. • Flexible Targeting: Apply NTP settings to specific devices or entire fabrics. • Persistent Configuration: NTP settings are stored in the XCO database for consistency across reboots and updates.

Table 2: Supported OS ONE features (continued)

Feature	Description
Network Essentials: Interface	Enables comprehensive visibility and management of network interfaces across devices in the XCO inventory as part of the Network Essentials module, which provides foundational services for fabric orchestration. The following features are supported: • Interface Listing: CLIs list interfaces on a device • Filter by interface Type: ◦ physical, loopback, ve, po, or all • State Monitoring: ◦ Admin State: up, down, or all ◦ Operational State: up, down, or all • Redundant Management Interface Listing: ◦ Option to list RM-enabled interfaces using --rmelist • Displays interface details: ◦ Device IP ◦ Interface name and type ◦ Admin and operational status ◦ Ensures that device configurations are synchronized with XCO and allows quick detection of drift or mis-configurations.
Support Maintenance Mode	Allows temporary restriction or modification to customer support tools and services during planned maintenance activities, system upgrades, or critical backend operations. This mode: • Ensures system stability during updates • Displays customizable maintenance messages to users • Automatically restores full support functionality post-maintenance • Logs maintenance activities for compliance and review Ideal for rolling out new support features or performing backend updates without disrupting live support.
Switch Firmware upgrade	Allows firmware upgrades on deployed fabrics (non-Clos, multi-rack, 3-Clos, 5-Clos) without traffic loss. Supports targeted upgrades using a prepared device list. Optimizes the upgrade process using Maintenance Mode (MM), Drift and Reconcile (DRC), and reboot options based on image build types.
Device and Network Level Settings	Introduces enhanced configuration capabilities at both the device and network levels: • Activates maintenance mode for device upgrades, initiates on-demand health checks, and configures single or periodic configuration backups. • Provides upcoming certificate and password expiry alerts. • Helps configure route load sharing maximum paths, unicast reverse path forwarding, MCT bring up delay.

Table 2: Supported OS ONE features (continued)

Feature	Description
Device Secure Settings	Enables security configurations on onboarded devices, whether part of a fabric or standalone in the XCO inventory. The following settings are supported: • Min-tls-version • Mac-algorithm • Key-exchange-algorithm • Cipher • Telnet • Max-password-age • Force-default-password-change
Device Password Expiry and Clear	Allows configuration of password expiry attributes on Extreme OS ONE devices. • Password expiry alerts are generated ahead of the expiration date, based on predefined thresholds. • Alert levels include: ◦ Info ◦ Minor ◦ Major ◦ Critical These levels help users take timely action to update passwords and maintain device security.
Inventory Switch LCM	Device certificates are installed and configured during OS ONE device registration in XCO.
Switch Health Management (SHM) and DRC	By default, health check functionality is disabled when Extreme OS ONE devices are registered. You can enable health checks on the device using the following operations: • **health-check-enable** • **health-check-interval** • **health-check-heartbeat-miss-threshold** XCO is configured to automatically back up device configurations every 6 minutes by default. XCO also supports Drift and Reconcile (DRC) at the device level. It compares the current device configuration with the expected configuration stored in XCO. If a drift is detected, XCO reconciles the configuration to restore the intended state. APIs are available to initiate drift and reconcile requests.
Switch System Support Save	Collects the system support save of the Inventory, Tenant, and Fabric service logs, and their associated database dumps.
Infrastructure	

Table 2: Supported OS ONE features (continued)

Feature	Description
TPVM Single Node	Deploys TPVM on OS ONE devices using IAH microservices. After deploying TPVM, you can install XCO using the same process as SLX-based TPVM deployments, with only minor differences in CLI commands. By default, XCO TPVM needs: • CPU: 2 cores • RAM: 4 GB • Disk: ◦ Rootfs – 10GB as per TPVM.img ◦ /apps – as per platform • 2 SSD devices: ◦ Available size 128 GB or 256 GB ◦ Default /apps size 50 GB • 1 SSD devices: ◦ Default /apps size – 30GB
TPVM Multinode	Deploys TPVM in a multinode configuration to achieve HA capabilities. To set up TPVM in multinode mode: 1. Deploy TPVM on two separate OS ONE devices. 2. Enable passwordless SSH between the two TPVM instances. This setup allows XCO to deploy required binaries and packages across both nodes without prompting for credentials.
TPVM Upgrade	Currently, only incremental TPVM upgrades are supported for the XCO 4.0.x release.
Notification Services	Sends notifications to external entities: • Events: Derived from syslog messages from managed devices. • Alerts: Triggered by unexpected conditions in XCO services. • Alarms: Stateful notifications raised and cleared by the system. • Tasks: User-initiated or scheduled operations (e.g., device registration, fabric creation). XCO 4.0.x OS ONE integration supports: • Alarms • Events • Fault Health Checks Supported Alerts and Alarms: • Port Flaps – Frequent interface up/down transitions. • Device Health – Monitors device operational status. • Certification Expiry – Notifies before certificates expire. • Password Expiry – Alerts for upcoming password expiration. • Fabric Health – Tracks network fabric status and integrity.

Table 2: Supported OS ONE features (continued)

Feature	Description
Simple Network Management Protocol (SNMP) Notifications	SNMP traps are alert messages sent from a remote SNMP-enabled device to a central collector, the SNMP Manager. Trap messages are the main form of communication between SNMP monitoring tools: • SNMP Agent • SNMP Manager XCO functions as the SNMP Manager for all EXTREME OS ONE devices and agents. Once a device is registered, XCO auto-configures it to send SNMP v3 traps. XCO receives traps from all devices in its inventory and acts as an SNMP proxy, forwarding SNMP v2 and v3 traps to an external trap receiver, if configured. Trap messages are the primary communication method between SNMP Agents and Managers, used to alert the SNMP Manager of events on remote devices.
Network Essentials: SNMP	XCO supports configuring the following SNMP components on EXTREME OS ONE devices: • SNMP Communities • SNMP Users • SNMP Host
SupportSave	The **efa system supportsave** script collects logs, database snapshots, pod logs, deployment details, and system support-save data. It compresses all collected information into a ZIP file for easy sharing with the Extreme support team during troubleshooting. The script captures: • Logs from XCO microservices • Host-level service logs • A snapshot of the database
Policy	
Compact Fabric Quality of Service (QoS) - Phase 1	XCO supports configuring QoS settings on Ethernet interfaces and Port Channels, including: • DSCP trust • PCP and DSCP ingress/egress maps Policy application is dynamic and adapts to the port role, whether part of a fabric, tenant, port channel, or tenant endpoint group.

Supported Platforms and Deployment Models for Fabric Skill

Support includes Server, Open Virtual Appliance (OVA), and TPVM deployment models, supported TPVM versions, supported SLX-OS software versions, and supported SLX devices.

Note

- OVA deployment model does not support HA.
- As a best practice, refer to the following Extreme validated support matrices for supported platforms and deployment models information.

Table 3: Server Deployment Models

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Server Version	Virtual Machine
3.4.x, 3.5.x, 3.6.x	More than 24	Yes	18.04 LTS and 20.04 LTS	• CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
3.7.x, 3.8.x	More than 24	Yes	20.04 LTS and 22.04 LTS	• CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
4.0.0	More than 24	Yes	20.04 LTS and 22.04 LTS	• CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
4.0.2	More than 24	Yes	22.04 LTS	• CPU: 4 cores • Storage: 64 GB • RAM: 8 GB

Table 4: OVA Deployment Models

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Virtual Machine
3.4.x, 3.5.x, 3.6.x	More than 24	Yes	20.04 LTS	• CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
3.7.x, 3.8.x	More than 24	Yes	22.04 LTS	• CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
4.0.x	More than 24	Yes	22.04 LTS	• CPU: 4 cores • Storage: 64 GB

Table 4: OVA Deployment Models (continued)

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Virtual Machine
				• RAM: 8 GB

Table 5: TPVM Deployment Models

XCO Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum Network OS Version
3.4.x, 3.5.x, 3.6.x	• SLX 9150 • SLX 9250 • SLX 9740 • Extreme 8520 • Extreme 8720 • Extreme 8820	Up to 24	Yes	20.04 LTS	SLX-OS 20.5.2a
3.7.x	• SLX 9150 • SLX 9250 • SLX 9740 • Extreme 8520 • Extreme 8720 • Extreme 8820	Up to 24	Yes	22.04 LTS	SLX-OS 20.6.3a
3.8.x	• SLX 9150 • SLX 9250 • SLX 9740 • Extreme 8520 • Extreme 8720 • Extreme 8820 (20.4.3 and later)	Up to 24	Yes	22.04 LTS	SLX-OS 20.7.1
4.0.0	• SLX 9150 • SLX 9250 • SLX 9740 • Extreme 8520 • Extreme 8720 • Extreme 8820 (20.4.3 and later)	Up to 24	Yes	22.04 LTS	SLX-OS 20.7.2
	Extreme 8730 (OS ONE	Up to 24	Yes	22.04 LTS	OS ONE 22.2.0.0

Table 5: TPVM Deployment Models (continued)

XCO Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum Network OS Version
4.0.2	• SLX 9150 • SLX 9250 • SLX 9740 • Extreme 8520 • Extreme 8720 • Extreme 8820	Up to 24	Yes	22.04 LTS	SLX-OS 20.7.3a
	• Extreme 8730 (OS ONE)	Up to 24	Yes	22.04 LTS	OS ONE 22.2.0.0, OS ONE 22.2.1.0, OS ONE 22.2.2.0

Table 6: TPVM Software Support

XCO Version	TPVM Version	Minimum Network OS Version
3.4.0	4.6.6	SLX-OS 20.5.3a
3.4.1	4.6.7	SLX-OS 20.5.3a
3.4.2	4.6.8	SLX-OS 20.5.3a
3.5.0	4.6.10	SLX-OS 20.6.1
3.6.0	4.6.13, 4.6.14	SLX-OS 20.6.2, SLX-OS 20.6.2a
3.7.0	4.6.17, 4.7.0	SLX-OS 20.6.3a
3.8.0	4.7.4	SLX-OS 20.7.1a
3.8.1	4.7.5	SLX-OS 20.7.1a
3.8.2	4.7.7	SLX-OS 20.7.1ab
4.0.0	4.7.7	SLX-OS 20.7.2
3.8.3	4.7.8	SLX-OS 20.7.1.ab
3.8.4	4.7.10	SLX-OS 20.7.2ab
3.8.5	4.7.12	SLX-OS 20.7.3a
3.8.6	4.7.13	SLX-OS 20.7.3a
3.8.7	4.7.14	SLX-OS 20.7.3a

Table 6: TPVM Software Support (continued)

XCO Version	TPVM Version	Minimum Network OS Version
3.8.8	4.7.15	SLX-OS 20.7.3ab
4.0.2	4.7.15	SLX-OS 20.8.1

Table 7: IP Fabric Topology Matrix

Device	SLX-OS/OS ONE Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9150	SLX-OS 20.6.x, 20.7.x, and 20.8.x	Yes	-	-	-	Yes
SLX 9250	SLX-OS 20.6.x, 20.7.x, and 20.8.x	Yes	Yes	Yes	-	Yes
SLX 9540	SLX-OS 20.6.x, 20.7.x, and 20.8.x	Yes	-	-	Yes	-
SLX 9640	SLX-OS 20.6.x, 20.7.x, and 20.8.x	-	-	-	Yes	-
SLX 9740	SLX-OS 20.6.x, 20.7.x, and 20.8.x	-	Yes	Yes	Yes	Yes
Extreme 8720	SLX-OS 20.6.x, 20.7.x, and 20.8.x	Yes	Yes	Yes	Yes	Yes
Extreme 8520	SLX-OS 20.6.x, 20.7.x, and 20.8.x	Yes	-	-	Yes	Yes
Extreme 8820	SLX-OS 20.6.x, 20.7.x, and 20.8.x	Yes	Yes	-	Yes	Yes
Extreme 8730	OS ONE 22.2.0.0, OS ONE 22.2.1.0, and OS ONE 22.2.2.0	-	-	-	-	Yes (2 Nodes)

XCO Upgrade Prerequisites

Prerequisites for XCO upgrade process with the default gateway changed:

1. Ensure that no DNS configuration exists under TPVM config and resolv.conf.
2. Presence of management connectivity from SLX and TPVM to external build server image, wherein image is available during SLX and TPVM upgrade process.

If file/etc/sshd/sshd_config is modified to non-default values, then manually readjust the following parameters:

- MaxStartups 30:30:100
- MaxAuthTries 6
- LoginGraceTime 120

Note
The hardening script, extr-granite.py bundled with EFA 2.6.1 will not automatically modify the above mentioned parameters.

Open Issues

The following issues are open in this release of the software.

Issue ID	Description
XCO-17877	**Symptom** Fabric health remains degraded after a leaf reboot in an MCT pair. In EFA, the BGP neighbor and MCT peer/keepalive states remain stale, even though the sessions have recovered on the device. **Condition** This issue occurs when raslog (syslog) events for BGP or MCT state changes are missed after a device reboot. In the absence of a subsequent raslog event, a full device update is not triggered, leaving the fabric state stale. **Workaround** Run the **`efa inventory device update`** command on the device to force a full device discovery and refresh all states. **Recovery** There is no automatic recovery. Fabric health returns to green only after the next full device discovery or by manually running `efa inventory device update` on the device.
XCO-17747	**Symptom:** When an EPG is created with the L3-handoff type and inventory device update is performed, a VE drift is observed for interface VE 8192(IRB VE) during drift detection. VE drift is observed on non-MCT devices when multiple L3-extension EPGs share the same VRF. **Condition:** This issue occurs during the following operations: • Create VRF with layer3-extension-enable true • Create EPG with L3-handoff type along with VRF and local-ip Perform inventory device update and identify the VE drift on that device for IRB VE(8192). In L3-extension EPG deployments where D1 and D2 are not an MCT pair, creating EPG1 on D1 (first EPG for the VRF) shows no drift, but creating EPG2 on D2 under the same VRF results in VE drift on D2.

Issue ID	Description
XCO-17656	**Symptom:** EFA commands may experience delayed responses during EFA service restart or reload; in some scenarios, continuous fabric service reboots are observed. **Condition:** This issue occurs when the fabric health calculation cron job and fabric configure run in parallel on a scaled fabric setup. Due to a timing conflict during health calculation, the issue may not occur consistently. **Recovery:** When fabric health is green and efa health is not green after the settings update, perform the following: 1. Update any of the fabric settings which does not cause any traffic flow hindrance: **`efa fabric setting update --name <fabric_name> --<setting_to_be_updated> <settings_value>`** Example: **`efa fabric setting update --name fab2 --max-paths 128`** 2. Configure Fabric: **`efa fabric configure --name <fabric_name>`** Example: **`efa fabric configure --name fab2`**
XCO-17545	**Symptom:** Auto/manual DRC fails after the default configuration is copied on a device that has a large number of route-maps configured. **Condition:** DRC (auto/manual) now runs after a copy default configuration or configuration backup and restore with default configuration on a device. **Recovery:** Restart the policy service.
XCO-17537	**Symptom:** After disabling EVPN neighbor activation (no activate) and reloading the device with MM enabled, the fabric health turns red and does not recover. The EFA fabric status shows `config_gen_reason` as `BGPU`. **Condition:** This issue occurs when: • EVPN neighbors are in deactivated state on the switch • Device reloaded with MM enabled **Recovery:** Run one of the following commands: • **`efa fabric debug device drift --force`** • **`efa fabric configure`**

Issue ID	Description
XCO-17485	**Symptom:** The SLX firmware downgrade fails on the BorderLeaf node with status "Firmware activation failure with auto recovery, XCO nodes are not in sync." **Condition:** The issue occurs during firmware download execution when a manually created group includes a TPVM device along with other devices. **Workaround:** Use the documented firmware download grouping: place devices hosting active and standby XCO nodes in separate groups, ensuring no other devices share the same group as the EFA-hosting device. For manual grouping (CLOS and non-CLOS), place TPVM1 in Group 1 and TPVM2 in Group 2.
XCO-17427	**Symptom:** During fabric configuration, the **MCT-Peer-KeepAlive-Enable** setting is enabled on the switch even though it is set to **disabled** at the fabric level when the fabric is in the **Fabric Created** state. **Condition:** This issue occurs under the following conditions in 3.8.5 and 4.0.2: • A fabric is created where mct-peer-keepalive-enable is enabled by default • Devices are added to the fabric • The fabric setting mct-peer-keepalive-enable is later changed to disabled • Fabric configuration is pushed to the switches Despite **mct-peer-keepalive-enable** being set to **disabled** at the fabric level, the switch receives and applies the setting as **enabled**. **Workaround:** Apply the MCT Peer Keepalive setting change either before adding devices to the fabric, or after the first successful fabric configure. If changes are needed post-configure, update the setting and then re-run fabric configure once the fabric reaches configure-success state.

Issue ID	Description
XCO-17371	**Symptom:** After a system restore, **`efa health show`** displays **RED**, while **`efa fabric health show`** displays **Green**. **Condition:** The **`efa system restore`** command overwrites the database using the backup. **Workaround:** Run system restore: " efa system restore" After the system restore completes, perform the following steps: 1. Clear monitor health: "efa health debug clear" 2. Clear alarm history: "efa system alert clear" 3. Recalculate health: "efa fabric debug health compute" **Recovery:** 1. Clear monitor health: "efa health debug clear" 2. Clear alarm history:"efa system alert clear" 3. Recalculate health:"efa fabric debug health compute" 4. Verify health status:"efa health show --detail" 5. If any irrelevant resource alarms are present, clear them: "efa system alarm close --resource <resourcepath> showing in efa health show"
XCO-17336	**Symptom:** Fabric health turns Black after evpn neighbour deactivate/activate. **Condition:** If the service restarts within 60 seconds after an EVPN neighbor deactivate/activate operation, BGP EVPN log events may not be processed, causing the fabric health to appear as Black. **Workaround:** Manually trigger a device update using "efa inventory device update --ip <device-ip>" for affected devices to refresh BGP neighbor states **Recovery:** The issue self-resolves after the next periodic timer-based discovery (typically 30–60 minutes) or when a configuration change triggers a full device update.

Issue ID	Description
XCO-17306	**Symptom:** After auto DRC, expected VRF config are not reflected in the switch. **Condition:** Tenant API Delete/Create PO, EPG or VRF with force option overlap with out of band updates of VRF. This is seen in the deep device discovery during the create/delete flow causing the inconsistencies. **Workaround:** Do not perform OOB configuration changes on devices within 15–20 seconds before or after Force Delete operations. **Recovery:** Manual recovery required: Delete and recreate affected VRFs/EPGs through Tenant API to resynchronize Tenant DB with actual device state.
XCO-17263	**Symptom:** Inventory Drift Reconcile fails while reconciling SNMP user with error `"netconf rpc [error] 'as a bad length/size"`. **Condition:** When an XCO-managed user is created using an unmanaged group, reconciliation of managed user may fail if a drift occurs, such as the unmanaged group being deleted from the device. **Workaround:** Delete and recreate the managed user in XCO..
XCO-17244	**Symptom:** Enabling or disabling the MCT peer keep-alive through XCO can cause a brief traffic drop during the configuration update. **Workaround:** Traffic loss is expected during configuration. **Recovery:** Traffic is expected to recover automatically after configuration completes.
XCO-16991	**Symptom:** After running `efa inventory admin-state down` followed by `admin-state up` on a spine device, the fabric health can turn Black and remain degraded for an extended period. During this time, XCO may report missing BGP sessions and physical links, and affected devices can remain in a configuration refresh error state even though all device-side BFD, BGP, and LLDP sessions are operational. **Recovery:** Fabric recovers automatically after the next cron job after 30 mins (max 60 mins). **Manual Recovery:** **`efa fabric debug health compute --name <fabric_name>`**

Issue ID	Description
XCO-16945	**Symptom:** When a local TPVM user (who is not a member of the `efa` group) logs into the EFA CLI directly, the security banner is not displayed. The banner is visible only when the `extreme` user is used for TPVM access. **Workaround:** Add the user to the `efa` user group.
XCO-16467	**Symptom:** DRC for prefix list/route-map does not work as expected. **Condition:** The issue occurs when OOB configuration contains case-duplicate prefix list names. **Workaround:** Do not use duplicate names (names that differ only by letter case) when creating OOB entries.
XCO-16017	**Symptom:** EFA commands may take a long time to respond during an EFA service restart or reload. In some cases, continuous fabric service reboots are observed. **Condition:** In scaled deployments with more than 15 devices, the fabric service starts before the tenant service is fully initialized. **Workaround:** Wait for the EFA commands to respond, then continue. • Up to ~4 hours in scaled setups with ~150 devices • Up to ~30 minutes in setups with ~25 devices This issue was observed once during the reboot timeframe. **Recovery:** Restart the fabric service after ensuring that the tenant service is already running.
XCO-15632	**Symptom:** When DRC is triggered simultaneously on all MCT devices after a system restore, DRC reconciliation can fail on one or more MCT devices during MCT cluster configuration reconciliation. **Workaround:** Execute DRC for the failed devices.

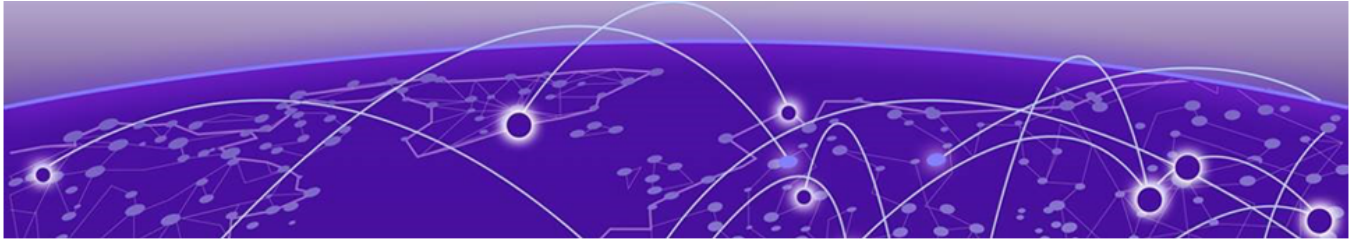

Help and Support

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

Extreme Portal

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

The Hub

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

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2800. For the support phone number in your country, visit www.extremenetworks.com/support/contact.

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

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

Subscribe to Product Announcements

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

1. Go to The Hub.
2. In the list of categories, expand the **Product Announcements** list.

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

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