



# ExtremeCloud™ Orchestrator Release Notes

Version 3.4.0

9037946-00 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/about-extreme-networks/company/legal/trademarks](https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks)

## Open Source Declarations

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

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



# Table of Contents

---

<b>Release Notes.....</b>	<b>4</b>
New In This Release.....	4
Supported Platforms and Deployment Models for Fabric Skill.....	6
Supported Platforms and Deployment Models for Visibility Skill.....	10
XCO Upgrade Prerequisites.....	12
Known Limitations.....	13
Known Limitations in Fabric Skill.....	13
Known Limitations in Visibility Skill.....	15
Defects Closed with Code Changes .....	16
Defects Closed without Code Changes.....	25
Open Defects.....	27
Help and Support.....	37
Subscribe to Product Announcements.....	37



# Release Notes

---

- [New In This Release](#) on page 4
- [Supported Platforms and Deployment Models for Fabric Skill](#) on page 6
- [Supported Platforms and Deployment Models for Visibility Skill](#) on page 10
- [XCO Upgrade Prerequisites](#) on page 12
- [Known Limitations](#) on page 13
- [Defects Closed with Code Changes](#) on page 16
- [Defects Closed without Code Changes](#) on page 25
- [Open Defects](#) on page 27
- [Help and Support](#) on page 37

## New In This Release

---

ExtremeCloud Orchestrator 3.4.0 introduces the following features and resolves a few issues through defect fixes. For information about XCO deployment, refer to the [ExtremeCloud Orchestrator Deployment Guide, 3.4.0](#).



### 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.

**Table 1: Features and Improvements**

Feature	Description
Tenant : BGP : Peer Group : Update Source : IPv6 Address	Learn how to Configure IPv6 Address as Update Source. For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a> .
Tenant : VRF : Default Information Originate	Learn how to enable or disable default information originate on a tenant VRF. For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a> .

**Table 1: Features and Improvements (continued)**

Feature	Description
Alarm Management	<p>Refer to the Alarm Inventory topic for PortFlap alarm details. Refer to the Health APIs for command to clear the health of a sub-tree for a specified resource.</p> <p>Refer to the Fabric Health Alerts topic for Port Flap alerts detail.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
QoS Profiling and Mapping at the Tenant and Port Level	<p>Learn how to configure a QoS Service Policy Map. Refer to the Tenant and Port Level Binding topic for XCO support of application of profile at tenant and port level.</p> <p>Refer to the Policy Service QoS Support topic with QoS service policy map details.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
SNMP traps for SLX device certificate expiry	<p>Learn how to configure the number of days before expiration of a certificate for an alert to be issued.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
Infra and Security Enhancements	<p>Show Device Adapter Connection Status describes how to show the connection status of device adapters.</p> <p>Show Device Certificate Expiry Time describes how to view the status of device certificates and their expiry time.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
Policy Service Enhancements	<p>Learn how to delete OOB entries from XCO database.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
XCO Licensing and Administration	<p>New topics "XCO License Service Management", "Software Licensing Overview", "Software Licensing Tasks", "Licensed Features and Part Numbers", "Licensing Supportsave Details", and "License Backup and Restore", describe the new XCO license service.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
Static IP support for dual stack IP type	<p>Updates to the "Configure Static IP Addresses for Management Sub Interfaces" topic with details on static IP support for Dual stack IP in an XCO deployment.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
Syslog certificate renewal	<p>Refer to the Syslog CA topic for certificate renewal procedure.</p> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator CLI Administration Guide, 3.4.0</a>.</p>
UI Enhancements	<p>The ExtremeCloud Orchestrator GUI has the following updates:</p> <ul style="list-style-type: none"> <li>Provision tenants using the UI</li> </ul> <p>For more information, refer to the <a href="#">ExtremeCloud Orchestrator GUI Administration Guide, 3.4.0</a>.</p>

**Table 1: Features and Improvements (continued)**

Feature	Description
Security enhancements	TACACS+ server configuration for XCO. For more information, refer to the <a href="#">ExtremeCloud Orchestrator Security Guide, 3.4.0</a> .
QoS profile options	Do not apply --service-policy and egress QoS maps together. For more information, refer to <a href="#">Known Limitations in Fabric Skill on page 13</a> .

For other additional information, see [Defects Closed with Code Changes](#) on page 16.

## 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 support platforms and deployment models information.

**Table 2: Server Deployment Models**

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Server Version	Virtual Machine
3.1.x	More than 24	Yes	16.04, 18.04, and 20.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.2.x	More than 24	Yes	18.04 and 20.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 2: Server Deployment Models (continued)**

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Server Version	Virtual Machine
3.3.x	More than 24	Yes	18.04 and 20.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.4.0	More than 24	Yes	18.04 and 20.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 3: OVA Deployment Models**

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Virtual Machine
3.1.x	More than 24	Yes	18.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.2.x	More than 24	Yes	18.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.3.x	More than 24	Yes	20.04 LTS	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.4.0	More than 24	Yes	20.04 LTS	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 4: TPVM Deployment Models**

XCO Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
3.2.x	<ul style="list-style-type: none"> <li>• SLX 9150</li> <li>• SLX 9250</li> <li>• SLX 9740</li> <li>• Extreme 8520</li> <li>• Extreme 8720</li> <li>• Extreme 8820 (20.4.3 and later)</li> </ul>	Up to 24	Yes	18.04	20.4.3
3.3.0	<ul style="list-style-type: none"> <li>• SLX 9150</li> <li>• SLX 9250</li> <li>• SLX 9740</li> <li>• Extreme 8520</li> </ul>	Up to 24	Yes	20.04 LTS	20.5.2

**Table 4: TPVM Deployment Models (continued)**

XCO Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
	<ul style="list-style-type: none"> <li>Extreme 8720</li> <li>Extreme 8820 (20.4.3 and later)</li> </ul>				
3.3.1	<ul style="list-style-type: none"> <li>SLX 9150</li> <li>SLX 9250</li> <li>SLX 9740</li> <li>Extreme 8520</li> <li>Extreme 8720</li> <li>Extreme 8820 (20.4.3 and later)</li> </ul>	Up to 24	Yes	20.04 LTS	20.5.2a
3.4.0	<ul style="list-style-type: none"> <li>SLX 9150</li> <li>SLX 9250</li> <li>SLX 9740</li> <li>Extreme 8520</li> <li>Extreme 8720</li> <li>Extreme 8820 (20.4.3 and later)</li> </ul>	Up to 24	Yes	20.04 LTS	20.5.2a

**Table 5: TPVM Software Support**

XCO Version	TPVM Version	SLX-OS Version
3.1.0	4.5.6	20.4.2a
3.1.1	4.5.8	20.4.3
3.2.0	4.5.10	20.4.3a
3.2.1	4.5.12	20.5.1
3.3.0	4.6.2	20.5.2
3.3.1	4.6.4	20.5.2a
3.4.0	4.6.6	20.5.3a

**Table 6: IP Fabric Topology Matrix**

Device	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9150	20.2.x, 20.3.x, 20.4.x	✓				✓
SLX 9250	20.2.x, 20.3.x, 20.4.x	✓	✓	✓		✓
SLX 9540	20.2.x, 20.3.x, 20.4.x	✓			✓	



**Table 6: IP Fabric Topology Matrix (continued)**

Device	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9640	20.2.x, 20.3.x, 20.4.x				✓	
SLX 9740	20.2.x, 20.3.x, 20.4.x		✓	✓	✓	✓
Extreme 8720	20.3.x, 20.4.x	✓	✓	✓	✓	✓
Extreme 8520	20.3.x, 20.4.x	✓			✓	✓
Extreme 8820	20.4.3		✓	✓	✓	✓

**Table 7: XCO or EFA, Neutron, and SLX-OS Compatibility**

XCO or EFA Version	Neutron Version	SLX-OS Version
2.5.4, 2.5.5	3.1.1-04	20.3.2d

## Supported Platforms and Deployment Models for Visibility Skill

Support includes Server, OVA, and supported devices and software.



### Note

- Upgrade from XVM (Extreme Visibility Manager) to XCO is not supported.
- XCO supports only a fixed set of special characters for hostnames. Any additional characters configured in MLX or SLX are reconciled in XCO and can be edited or deleted. Any configuration name must start with an alphanumeric character and can contain a-z A-Z 0-9 \_ -.

**Table 8: Ubuntu Server Version**

XCO Version	Ubuntu Version	Virtual Machine
3.1.x	18.04 and 20.04	Minimum: <ul style="list-style-type: none"><li>• CPU: 4 cores</li><li>• Storage: 128 GB</li><li>• RAM: 8 GB</li></ul> Recommended: <ul style="list-style-type: none"><li>• CPU: 16 cores</li><li>• Storage: 200 GB</li><li>• RAM: 32 GB</li></ul>
3.2.x	18.04 and 20.04	Minimum: <ul style="list-style-type: none"><li>• CPU: 4 cores</li><li>• Storage: 128 GB</li><li>• RAM: 8 GB</li></ul> Recommended: <ul style="list-style-type: none"><li>• CPU: 16 cores</li><li>• Storage: 200 GB</li><li>• RAM: 32 GB</li></ul>

**Table 8: Ubuntu Server Version (continued)**

XCO Version	Ubuntu Version	Virtual Machine
3.3.x	18.04 and 20.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 128 GB</li> <li>• RAM: 8 GB</li> </ul> Recommended: <ul style="list-style-type: none"> <li>• CPU: 16 cores</li> <li>• Storage: 200 GB</li> <li>• RAM: 32 GB</li> </ul>
3.4.0	18.04 and 20.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 128 GB</li> <li>• RAM: 8 GB</li> </ul> Recommended: <ul style="list-style-type: none"> <li>• CPU: 16 cores</li> <li>• Storage: 200 GB</li> <li>• RAM: 32 GB</li> </ul>

**Table 9: OVA Deployment Models**

XCO Version	Ubuntu Version	Virtual Machine
3.1.x	18.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.2.x	18.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 9: OVA Deployment Models (continued)**

XCO Version	Ubuntu Version	Virtual Machine
3.3.x	20.04 LTS	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.4.0	20.04 LTS	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 10: Supported Devices and Software**

Device	Supported Software
Extreme 9920	Extreme 9920 software with the NPB application <ul style="list-style-type: none"> <li>• 21.1.2.x</li> </ul>
Extreme Routing MLX Series	<ul style="list-style-type: none"> <li>• NetIron 6.3.00 patches</li> </ul>
Extreme Switching SLX 9140	<ul style="list-style-type: none"> <li>• SLX-OS 18s.1.03 patches</li> </ul>
Extreme Switching SLX 9240	<ul style="list-style-type: none"> <li>• SLX-OS 18s.1.03 patches</li> </ul>

## 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.

## Known Limitations

---

Note the following caveat for this release of ExtremeCloud Orchestrator.

- If the CLOS setup firmware upgrade encounters the error "Cannot start download before the new image is committed", then create a separate group only for the active XCO node and perform the firmware upgrade.

## Known Limitations in Fabric Skill

Follow these caveats and limitations when using Fabric Skill.

*Quality of Service (QoS) policy service support*

- The XCO-driven application of policy is dynamic and can vary depending on the port's role, whether it belongs to a fabric, tenant, port channel, or tenant endpoint group.



### Tip

As a best practice, avoid running user-driven policy operations in parallel with fabric, tenant, port channel, and tenant endpoint group operations.

To ensure that the fabric, tenant, port channel, and tenant endpoint group configurations are effective, run the **show** command before proceeding with the policy operations, and vice-versa.

- Before running the force operations, including deletion, ensure that you unbind the policies (QoS) from all the relevant targets (fabric, tenant, port, port channel, and tenant endpoint group) to avoid stale policies (QoS) in the system.
- Before executing the QoS policy bind commands, remove any conflicting or additional OOB (Out of Band) QoS configurations from the switches to ensure that the correct policies are applied to the ports.
- There is no support for a lossless hardware profile. Therefore, you must switch the configuration on SLX devices to a lossy hardware profile before provisioning QoS policies from XCO.
- There is no support for egress QoS maps. While XCO allows the configuration of egress QoS maps, as a best practice, do not configure any egress QoS maps from XCO due to limitations in SLX support of egress QoS maps.

*VRF delete from EPG and re-adding VRF to EPG fails intermittently*

Symptom	Condition	Workaround
Endpoint group (EPG) update <b>vrf-add</b> operation fails with the reason as VRF to be added has conflicting VRF on the switch.	Run EPG update <b>vrf-add</b> , <b>vrf-delete</b> , and <b>vrf-add</b> operation CLI in quick succession: <ol style="list-style-type: none"> <li>1. Update EPG for operation <b>vrf-add</b>.</li> <li>2. Update EPG for operation <b>vrf-delete</b>.</li> <li>3. Update the same EPG again with operation <b>vrf-add</b> for the same VRF which was deleted in step 2.</li> </ol>	Wait of 30 seconds between the EPG update <b>vrf-add</b> and vrf-delete operations on the same EPG.

*REST operations are not retried (as applicable) during the service boot*

Symptom	Condition	Workaround
REST operations are not retried (as applicable) during the service boot up.	After publishing the necessary events on the message bus, the status for the REST operations are not set automatically.	Manually set the status for all REST operations.

*RBAC: XCO shows "export EFA\_TOKEN" command suggestion when a tenant user logs in*

Symptom	Condition	Workaround
XCO shows a <b>export EFA_TOKEN</b> message after a tenant user with RBAC logs in to the system.	When a user is created with the default login shell as sh.	XCO supports only bash shell for login or any other CLI commands. Ensure that bash is specified as the default login shell for all XCO user accounts.

**EFA Token command message**

Please type this in your shell:

**export**

```
EFA_TOKEN=eyJhbGciOiJSUzI1NiIsImtpZCI6IjEuMCI6IjE5cCI6IkpXVCJ9.eyJjb21tb25fbmFtZSI6IkgVQSBUB2t1b1BTZXJ2aWN1IiwidWFzIjpbeyJ0YXJnZXQiOiJFRkEiLCJyb2x1IjoilIyLVRudEFkbWluInldLCJvcml0iJFeHRyZW11IE5ldHdvcmtzIiwidmVyIjoimS4wIiwiaWQiOiIiLCJleHAiOjE2NDUyNDcxNDIsImp0aSI6IjZjMjA4ZDUxLTkwNzgtMTFlYy1lZjk5LWNhZk1kMDY1YzIwNyIsImh0dCI6MTY0NTE2Mdc0MiwiaXNzIjoiaRUZBIFRva2VuIFNlcnZpY2UiLCJyYmYiOjE2NDUxNjA3NDIsInN1YiI6InVzZXIyIn0.b7m5PINijeEdNSqntE2ZUrqKLKQAU079vXyBI dgHbXKt9ULfa03vMU1j fBO1qFb1-
```

```
x0oHmsAQ0pSsF5JLeMaMzMf1Lf78ktZ08U5IePq72vM5en35IR-
DNLyoGIZBeFeG6ZbBMoETzz5vf9OuefgQID3YdjcaLr7y1lCgDmLVFlgson77yCBpkTK15xm
1GRbtL7JKXZzShBE7E3kdW7N71MdM85Gc3r41-c8sfz7eo06gKrfTq9wXCv4_LVzR6-
KRSg6NyLq363WEpcK1A2Hs0Wo3T9TpquYHNaCWA5I1QTsG-
RHFdg4kxZP2fQpUp6Bgy1s6k59PVPn4-M-a81A- Time Elapsed: 4.619465187s —
```

*XCO CLI or REST request with scale config takes longer than 15 minutes fails*

Symptom	Condition	Workaround
<p>Tenant2 delete is successful whereas deleting Tenant1 took more than 15 minutes and failed with the following message:            Error : service is not available or internal server error has occurred, please try again later            Tenant service was running.            Tenant1 was not available after the error.</p>	<p>When you try to delete tenants in a single rack small data center deployment configured with scale tenant config</p>	<p>Any CLI or REST tenant operations, and any fabric operations taking more than 15 minutes, will timeout at the client side. The operation completes in the background. Run the <b>efa tenant show</b> command to view the actual state of the operation.</p>

## Known Limitations in Visibility Skill

Follow these caveats and limitations when using the Visibility Skill.

*LAG created when port channel deployment fails*

Any changes to ExtremeCloud Orchestrator configuration are reverted when a port channel deployment fails. However, a link aggregation group (LAG) is created on the device. The LAG is immediately deleted, but you can see the creation and deletion of a LAG in the device logs.

*MLX UDA profile must be associated with an ingress group if the policy contains a UDA match*

(MLX only) When you create an ingress group and associate it with an ingress policy, you must also associate the group with a UDA profile if the ingress policy contains a UDA match. For more information, see [ExtremeCloud Orchestrator GUI Admin Guide](#).

*Firmware upgrade requires an absolute path to image locations*

In the **Absolute Path** field, enter the complete file path to the location of the firmware image. The following are sample file paths for the various supported devices.

- Extreme 9920 (absolute path to the binary file): /root/TierraOS--NPB.bin
- SLX (absolute directory path where supported image files are located): /root/slxos18s.1.03/slxos18s.1.03a
- MLX (path to the manifest file): XMR-MLX/MLX\_npb\_06200\_mnf.txt

For more information, see [ExtremeCloud Orchestrator GUI Admin Guide](#).

*Device discovery*

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

*Listener policy byte count is incorrect when truncation is enabled*

On the Extreme 9920 device, the byte count for truncated packets is the actual byte count seen by the egress ACL before truncation.

## Defects Closed with Code Changes

---

The following defects were resolved in ExtremeCloud Orchestrator 3.4.0.

<b>Parent Defect ID:</b>	XCO-4127	<b>Issue ID:</b>	XCO-4127
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 3.0.0
<b>Symptom:</b>	Ports are not listed in the port-channel creation for SLX NPB devices		
<b>Condition:</b>	Even though the ports are not used in any other configurations, the ports are not listed in the port-channel creation. For these ports, speed is set to auto-negotiation, and ports are not connected with cable.		
<b>Workaround:</b>	For breakout ports, make sure that cables are connected so that port speed will be updated.		
<b>Recovery:</b>	N/A		

<b>Parent Defect ID:</b>	XCO-4146	<b>Issue ID:</b>	XCO-4146
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 2.7.2
<b>Symptom:</b>	The fabric devices continue to remain in cfg-refresh-err state after the tpvm fail over.		



<b>Parent Defect ID:</b>	XCO-4146	<b>Issue ID:</b>	XCO-4146
<b>Condition:</b>	<p>1.Fabric devices are already in cfg-refresh-err state due to LLDP Link down(LD) event.</p> <p>2. Bring up the LLDP links responsible for the fabric devices to be in cfg-refresh-err state.</p> <p>3. Execute the TPVM failover by 'tpvm stop' and 'tpvm start' commands during the LLDP Link up (LA) event handling caused by 2.</p>		
<b>Recovery:</b>	<p>1The user triggers LD/LA event by flapping the interface links which are the devices are in the cfg-refreshed state even though DRC wouldn't help out to recover the device to the cfg-sync state and the pending reason is "LA/LD".</p> <p>1.1. "shutdown" the interface link on the physical link on Devices follow by "efa inventory device update --ip &lt;device-ip&gt;", which generates LD events</p> <p>21.. "no shutdown" the interface link on the physical link on Devices follow by "efa inventory device update --ip &lt;device-ip&gt;", which generates LA events</p> <p>1.3. If the pending config contains "LA" : Execute "efa inventory drift-reconcile execute --ip &lt;device-ip&gt; --reconcile" on the devices which are in cfg-refresh-err /cfg-refreshed state [or] IF the pending config contains "LD,LA" : Execute "efa fabric configure --name &lt;fabric-name&gt;" to clean up the configuration on devices which are in cfg-refresh-err /cfg-refreshed state.</p> <p>[OR]</p> <p>2. The user reboots the devices without maintenance mode which are the devices are in cfg-refreshed state even though DRC wouldn't help out to recover the device to the cfg-sync state.</p> <p>2.1. "reload" the switches without out maintenance mode to enable</p> <p>2.2. Execute "efa inventory drift-reconcile execute --ip &lt;device-ip&gt; --reconcile" on the devices which are in cfg-refresh-err /cfg-refreshed state.</p>		

<b>Parent Defect ID:</b>	XCO-7183	<b>Issue ID:</b>	XCO-7183
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.2
<b>Symptom:</b>	<p>After changing DNS nameservers in /etc/netplan and running the update-dns.sh --dns-action allow, the following error is seen:</p> <pre>(efa:ubuntu)ubuntu@efa:/opt/efa\$ sudo /opt/efa/update-dns.sh /opt/efa/update-dns.sh Usage: --help - Show this message --dns-action &lt;'allow' 'disallow'&gt; - Allow host DNS entries to be forwarded to the pods (efa:ubuntu)ubuntu@efa:/opt/efa\$ sudo /opt/efa/update-dns.sh --dnsaction allow Unexpected nameserver entry of 127.0.0.53 found in /etc/resolve.conf (efa:ubuntu)ubuntu@efa:/opt/efa\$</pre>		

Parent Defect ID:	XCO-7183	Issue ID:	XCO-7183
<b>Condition:</b>	<p>In 18.04.6 and 20.04, Ubuntu uses a stub-resolv.conf located in /run/systemd/resolve/stub-resolv.conf . This file is symlink to /etc/resolv.conf in /run/systemd/resolve/.</p> <p>There is another file, resolv.conf which contains the information for DNS from netplan.</p> <p>Additionally, systemd-resolved provides a local DNS stub listener on IP address 127.0.0.53 on the local loopback interface. Programs issuing DNS requests directly, bypassing any local API may be directed to this stub, in order to connect them to systemd-resolved.</p> <p>Note: The best practice is for local programs to use the glibc NSS or bus APIs instead (as described above), as various network resolution concepts (such as link-local addressing, or LLMNR Unicode domains) cannot be mapped to the unicast DNS protocol.</p> <p>We do not recognize the 127.0.0.53 address as valid.</p>		
<b>Workaround:</b>	<p>If updating DNS to allow host entries to be forwarded to the pods using the update-dns.sh script in XCO-3.3.0 on Ubuntu 20.0.4 or 18.0.4-6 or above, follow these steps.</p> <p>20.0.4 or 18.0.4-6 or above, follow these steps:</p> <p>After netplan is applied and before running update_dns.sh</p> <ol style="list-style-type: none"> <li>1. Check if symlink exists, if not directly edit /etc/resolv.conf to netplan ip  <pre>ip \$ ls -l /etc/resolv.conf lrwxrwxrwx 1 root root 39 Feb 20 2021 /etc/resolv.conf -&gt; ../run/systemd/resolve/stub-resolv.conf &lt;&lt;&lt;symlink exists sbr@sbr-virtual-machine ~ \$</pre> </li> <li>2. Check if it has 127.0.0.53 ip in below files:  <pre>~ \$ cat /etc/resolv.conf   grep nameserver nameserver 127.0.0.53 sbr@sbr-virtual-machine ~ \$ cat /run/systemd/resolve/stub-resolv.conf   grep nameserver nameserver 127.0.0.53 sbr@sbr-virtual-machine ~ \$</pre> </li> <li>3. Edit the following file to add netplan ip for the nameserver and remove 127.0.0.53  <pre>sudo vi /run/systemd/resolve/stub-resolv.conf</pre> </li> <li>4. Check if both files are updated.  <pre>~ \$ cat /run/systemd/resolve/stub-resolv.conf   grep nameserver nameserver 10.10.10.0 sbr@sbr-virtual-machine ~ \$ cat /etc/resolv.conf   grep nameserver nameserver 10.10.10.0</pre> </li> </ol>		

<b>Parent Defect ID:</b>	XCO-7183	<b>Issue ID:</b>	XCO-7183
	sbr@sbr-virtual-machine ~ \$ 5. Run update_dns.sh --dns-action allow. 6. Run sudo netplan apply to restore /etc/resolv.conf and /run/systemd/resolve/stub-resolv.conf to its default value of 127.0.0.53.		

<b>Parent Defect ID:</b>	XCO-7955	<b>Issue ID:</b>	XCO-7955
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	When triggering the "Firmware Activate" process, it can lead to either parallel or serial execution, irrespective of the behavior of grouping devices for traffic loss. In cases where auto-commit is enabled, the activation can result in a "Firmware Commit Failed" status on the EFA end, even though the firmware commit has been successfully completed on the device end.		
<b>Condition:</b>	The "Firmware Activate" process is initiated from the user interface, either through the Inventory Page or the Fabric-wide Page, even in the midst of an incomplete operation on a subset of devices. For instance: Device 1 and Device 2 trigger a download with auto-commit enabled from either the Inventory or Fabric-wide Page. Device 3 triggers a download from the Fabric or Inventory Page. Subsequently, Device 1 and Device 2 attempt to continue with the "Activate Download" operation from the inventory or fabric page, resulting in a "Firmware Commit Failed" failure.		
<b>Workaround:</b>	Do not initiate firmware upgrades on other devices until the device completes both the Activate operation and the commit operation.		
<b>Recovery:</b>	Based on the error in the flow sequences, use the following set of commands: "efa inventory debug unblock-from-fwdl" , "efa inventory device firmware-download" to continue with download operation		

<b>Parent Defect ID:</b>	XCO-8070	<b>Issue ID:</b>	XCO-8070
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	'efa system backup' fails with Error : Failed to execute service lock API due to error Role ServiceAdmin does not have permissions to path: /v1/inventory/lockservice, method: POST.		
<b>Condition:</b>	When db error occurs in Rbac during initialization, then return gracefully. Allowing Rbac to do re-initialization. Earlier we were not returning when Db is unavailable/down, hence it was failing to load policies.		

Parent Defect ID:	XCO-8070	Issue ID:	XCO-8070
	Thus, not able to resolve permissions to run command.		

Parent Defect ID:	XCO-8128	Issue ID:	XCO-8128
Product:	XCO	Reported in Release:	XCO 3.2.1
Symptom:	Unable to Create EPG when out of band static-routes with next-hop interface exist		
Condition:	<p>Root cause:            There are multiple issues observed when oob static routes are added. (a) When null route is added, inventory table is populated with wrong interface name - tengigabitethernet 0, (b) published vrfupdate event does not contain interface name and nexthoptype value is set to invalid value - 0, (c) when tenant reads the event it writes incomplete information in its db and (d) Tenant does not know how to read back the incomplete back properly.</p>		
Workaround:	Not appilcable.		
Recovery:	<ol style="list-style-type: none"> <li>1. Remove the static-routes from the vrf config in the device.</li> <li>2. Update the tenant db vrf_static_route to delete the corresponding stale records (oob_created=1 and nh_type = 0).</li> <li>3. Create the previously failing new epg-s and confirm that they succeed.</li> <li>4. Re-add the static-routes in the vrf config in the device.</li> </ol>		

Parent Defect ID:	XCO-8200	Issue ID:	XCO-8200
Product:	XCO	Reported in Release:	XCO 3.3.0
Symptom:	SLX Devices are not allowed to execute the same firmware download execution flow, which could result in traffic loss. For example, it is not allowed to choose two Leaf devices from the same MCT pair.		
Condition:	From the User Interface, go to the <b>Fabric</b> page & select a few devices. Go to table action and select <b>Firmware Upgrade</b> option.		
Workaround:	The user selects the left-side leaf of the MCT pair and triggers firmware download and activation. Similarly, the user selects the right-side leaf of the MCT pair and triggers firmware download and activation.		
Recovery:	Choose another set of devices that will not result in traffic loss and proceed with the firmware download operation.		

Parent Defect ID:	XCO-8232	Issue ID:	XCO-8232
Product:	XCO	Reported in Release:	XCO 3.3.0
Symptom:	Error is observed while updating EFA system CLI setting		

<b>Parent Defect ID:</b>	XCO-8232	<b>Issue ID:</b>	XCO-8232
	Error : error creating directory on remote: Could not chdir to home directory /users/home21/<username>: No such file or directory		
<b>Condition:</b>	While using CLI "efa system settings update --remote-server-ip <ip> --remote-transfer-protocol scp --remote-server-username <username> --remote-server-password <password> --remote-server-directory <remote-server-directory>"		
<b>Workaround:</b>	Use Remote Server which has bash support installed.		
<b>Recovery:</b>	Add bash support and retry the CLI command.		

<b>Parent Defect ID:</b>	XCO-8234	<b>Issue ID:</b>	XCO-8234
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	The fabric alarm and the alarm status update notifications can briefly reflect a small time window where the fabric alarm is cleared when it is actually unhealthy.		
<b>Condition:</b>	This can occur during fabric formation or during any operation where fabric health is degraded due to multiple reasons (example:- spine to leaf link going down, BGP neighborhood going down between spine and leaf, etc...). Once a specific device and links are repaired and deemed healthy, the overall fabric alarm may temporarily be cleared although other devices remain unhealthy. Then subsequently the fabric alarm will be corrected and put into an unhealthy state due to the remaining unhealthy devices.		
<b>Workaround:</b>	N/A		
<b>Recovery:</b>	The fabric alarm automatically recovers to the proper state. Its just that the fabric alarm may temporarily be cleared when it is actually not cleared yet.		

<b>Parent Defect ID:</b>	XCO-8267	<b>Issue ID:</b>	XCO-8267
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	Devices successfully finished their backups and took 4 minutes but the completion timeout was set to 3 minutes. Hence we observed "Config Backup timed out for Device" error message.		
<b>Condition:</b>	The monitor process is making rest calls to Inventory every 2 seconds to see if the backup is done yet and after 3 minutes monitor claims failure. Setting the completion timeout greater than the netconf timeout which is 4 minutes and 50 seconds, so that the monitor won't have false positive failure messages.		

<b>Parent Defect ID:</b>	XCO-8289	<b>Issue ID:</b>	XCO-8289
	XCO	<b>Reported in Release:</b>	XCO 3.2.1

<b>Parent Defect ID:</b>	XCO-8289	<b>Issue ID:</b>	XCO-8289
<b>Product:</b>			
<b>Symptom:</b>	After setting breakout port on 9740, the next ports are still shown in 'show-running-config'		
<b>Condition:</b>	On 9740 breakout of port, can see unacceptable lines in efa show-running-config (with 21 broken out, can't have 22 admin state)		

<b>Parent Defect ID:</b>	XCO-8366	<b>Issue ID:</b>	XCO-8366
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.1
<b>Symptom:</b>	IPv6-Prefix over IPv4-Peer device setting under Inventory service becomes refreshed and gets removed from the device when device is removed from fabric or entire fabric gets deleted. This setting doesn't get applied automatically to the device when it is added back to the fabric or fabric is reconfigured.		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1. Configure fabric.</li> <li>2. Enable IPv6-Prefix over IPv4-Peer device setting from inventory CLI.</li> <li>3. Remove device from fabric or delete entire fabric.</li> <li>4. Add device back in fabric or re-configure fabric.</li> </ol> <p>Performing Step 4 doesn't configure IPv6-Prefix over IPv4-Peer setting on device and Inventory service keep on identifying drift for the same.</p>		
<b>Recovery:</b>	Run DRC from Inventory service before/after adding device to fabric & reconfiguring fabric		

<b>Parent Defect ID:</b>	XCO-8574	<b>Issue ID:</b>	XCO-8574
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	Delete/Remove route-map was successful even when bindings associated with BGP neighbor. It supposed to deny.		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1. Create route-map stanza</li> <li>2. Configure it on the device</li> <li>3. Create BGP peer and peer-group with route-map binding</li> <li>4. Delete the route-map with --seq all</li> </ol>		
<b>Workaround:</b>	Remove BGP peer/peer-group association first and then delete/remove the route-map from Device.		
<b>Recovery:</b>	Re-add the route-map to device again and then follow the workaround above for proper removal.		

<b>Parent Defect ID:</b>	XCO-8698	<b>Issue ID:</b>	XCO-8698
	XCO	<b>Reported in Release:</b>	EFA 2.7.2

<b>Parent Defect ID:</b>	XCO-8698	<b>Issue ID:</b>	XCO-8698
<b>Product:</b>			
<b>Symptom:</b>	Some of the anycast IP configs of existing ports-ctags of an epg have been found to be marked as 'deleted' in XCO's database. Could've been due to some sync issue that existed in 2.7 or previous images. When a port-group-add is performed on this epg with ports from new devices, the command is failed with 503 - Service available error as the tenant service gets a panic. Re-execution of same command succeeds without provisioning anycast configs in the device.		
<b>Condition:</b>	<p>Data corruption:</p> <p>There is no mapping of anycast-ip with VE port in endpoint_group_network_properties_ip table for select VE ports (3 out of 20 to be exact). Instead, a NULL value was seen in device_id and device_port_ip_id column. This is normally done only when port-group-delete is executed, not otherwise.</p> <p>Root cause:</p> <p>there are other associated epg port's anycast configs marked as deleted in XCO's database. This causes XCO to attempt to prepare configuration for these ports as well, as part of port-gruop-add use case. This use case has a bug in Tenant software and causes a non-fatal panic. this has cascaded to the actual issue customer faced.</p>		
<b>Workaround:</b>	Not applicable		
<b>Recovery:</b>	Not applicable		

<b>Parent Defect ID:</b>	XCO-8700	<b>Issue ID:</b>	XCO-8700
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	GUI upgrade went successful, but device stuck on GUI so can't perform configuration or delete the device.		
<b>Condition:</b>	GUI upgrade went successful, but device stuck on GUI so can't perform configuration or delete the device.		

<b>Parent Defect ID:</b>	XCO-8827	<b>Issue ID:</b>	XCO-8827
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	"efa system backup --remote" failed when password length exceeds 16 chars on the remote server.		
<b>Condition:</b>	System backup was failing due to error in decrypting the password. Hence unable to do scp to remote host.		

<b>Parent Defect ID:</b>	XCO-8831	<b>Issue ID:</b>	XCO-8831
	XCO	<b>Reported in Release:</b>	XCO 3.3.0

<b>Parent Defect ID:</b>	XCO-8831	<b>Issue ID:</b>	XCO-8831
<b>Product:</b>			
<b>Symptom:</b>	XCO Visibility showing non-existent ports on 9920.		
<b>Condition:</b>	When looking at ports on 9920 non-existent ports shows up sometimes.		

<b>Parent Defect ID:</b>	XCO-8935	<b>Issue ID:</b>	XCO-8935
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	Delete/Remove route-map from Device was successful even when bindings associated with BGP neighbor. It supposed to deny.		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1. Create route-map stanza.</li> <li>2. Configure it on the device.</li> <li>3. Create BGP peer and peer-group with route-map binding.</li> <li>4. Update route-map with operation remove-device.</li> </ol>		
<b>Workaround:</b>	Remove BGP peer/peer-group association first and then delete/remove the route-map from Device		
<b>Recovery:</b>	Re-add the route-map to device again and then follow the workaround above for proper removal.		

<b>Parent Defect ID:</b>	XCO-8936	<b>Issue ID:</b>	XCO-8936
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	Delete or remove prefix-list from Device was successful even when bindings associated with BGP neighbor. It supposed to deny.		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1. Create prefix-list.</li> <li>2. Create route-map.</li> <li>3. Create route-map-match and associate prefix created in Step 1 with route-map created in Step 2.</li> <li>4. Associate/advertise prefix-list/route-map in bgp peer-group.</li> <li>5. Delete or Remove Prefix-list from Device created in Step 1.</li> </ol>		
<b>Workaround:</b>	Remove BGP peer/peer-group association first and then delete/remove the prefix-list from Device.		
<b>Recovery:</b>	Re-add the prefix-list to device again and then follow the workaround above for proper removal.		

<b>Parent Defect ID:</b>	XCO-8975	<b>Issue ID:</b>	XCO-8975
	XCO	<b>Reported in Release:</b>	XCO 3.3.1



<b>Parent Defect ID:</b>	XCO-8975	<b>Issue ID:</b>	XCO-8975
<b>Product:</b>			
<b>Symptom:</b>	ICL Expansion - Fabric Event 'Link Add (LA)' Not Shown for All Switches.		
<b>Condition:</b>	LA should come on both device d1, d2 (or) if this is not feasible, then the PR changes solves the issue by having additional checks on fabric LA handling specific to mct devices.		

<b>Parent Defect ID:</b>	XCO-9323	<b>Issue ID:</b>	XCO-9323
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	GUI for Device Ports status when Group by "Name" doesn't match		
<b>Condition:</b>	GUI for Device Ports status when Group by "Name" doesn't match		

## Defects Closed without Code Changes

The following defects are closed without code changes in ExtremeCloud Orchestrator 3.4.0.

<b>Parent Defect ID:</b>	XCO-5263	<b>Issue ID:</b>	XCO-5263
<b>Reason Code:</b>	Already Reported		
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 3.1.0
<b>Symptom:</b>	Failed to report telemetry not streamed from device		
<b>Condition:</b>	When SLX device is discovered from XCO and statistics are getting streamed using telemetry service. Device is not sending statistics using telemetry service.		
<b>Workaround:</b>	Select the individual device and verify the statistics are streaming from device.		

<b>Parent Defect ID:</b>	XCO-8072	<b>Issue ID:</b>	XCO-8072
<b>Reason Code:</b>	Insufficient Information		
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	When configuring an OOB QoS map of traffic-class-cos type without DP value provided on SLX, after device update, the entry showed up on EFA with DP value of 4.		
<b>Condition:</b>	When configuring an OOB QoS map of traffic-class-cos type without DP value provided on SLX, after device update, the entry showed up on EFA with DP value of 4.		

<b>Parent Defect ID:</b>	XCO-8072	<b>Issue ID:</b>	XCO-8072
<b>Workaround:</b>	No workaround for this OOB entry.		
<b>Recovery:</b>	Delete this OOB entry from SLX device side.		

<b>Parent Defect ID:</b>	XCO-8131	<b>Issue ID:</b>	XCO-8131
<b>Reason Code:</b>	Not a Software Defect		
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.0
<b>Symptom:</b>	RMA command giving the error in EFA		
<b>Condition:</b>	RMA command giving the error in EFA		
<b>Workaround:</b>	After replacing with new device use same port pairs connection which were used before replacement and also If ports are in shutdown state, bring the ports up and re-run RMA.		
<b>Recovery:</b>	After replacing with new device use same port pairs connection which were used before replacement and also If ports are in shutdown state, bring the ports up and re-run RMA.		

<b>Parent Defect ID:</b>	XCO-8403	<b>Issue ID:</b>	XCO-8403
<b>Reason Code:</b>	Already Reported		
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	When adding SLX devices to fabric with add-bulk it fails		
<b>Condition:</b>	Add SLX 8820 device(while keeping all ports admin down) to fabric using add-bulk command and see if it fails		
<b>Workaround:</b>	Upgrade to 3.3.0 or above to have "auto admin enable feature" support for SLX 8820 device or enable all ports admin up and the try to add to fabric using add-bulk		
<b>Recovery:</b>	Upgrade to 3.3.0 or above to have "auto admin enable feature" support for SLX 8820 device or enable all ports admin up and the try to add to fabric using add-bulk		

<b>Parent Defect ID:</b>	XCO-8423	<b>Issue ID:</b>	XCO-8423
<b>Reason Code:</b>	Not Applicable		
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	Disk high usage and EFA is down on standby TPVM		
<b>Workaround:</b>	efactl clean commands is used to clean up the disk.		
<b>Recovery:</b>	efactl clean commands is used to clean up the disk.		

## Open Defects

The following defects are open in ExtremeCloud Orchestrator 3.4.0.

<b>Parent Defect ID:</b>	XCO-3445	<b>Issue ID:</b>	XCO-3445
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 3.0.0
<b>Symptom:</b>	DRC will not identify the drift and hence will not reconcile the drifted configuration		
<b>Condition:</b>	<p>Below are the steps to reproduce the issue:</p> <ol style="list-style-type: none"> <li>1. Configure multi rack Non-CLOS fabric.</li> <li>2. Manually remove the below set of configurations on device under:  <pre>router-bgp no neighbor 172.x.x.x password xxxx no neighbor 172.x.x.x update-source loopback 1 no neighbor 172.x.x.x peer-group overlay-ebgp-group address-family l2vpn evpn no retain route-target all</pre> </li> <li>3. Run "efa inventory drift-reconcile execute --ip &lt;device-ip&gt;"</li> </ol>		
<b>Recovery:</b>	Manually reconfigure the removed configurations from the device		

<b>Parent Defect ID:</b>	XCO-3471	<b>Issue ID:</b>	XCO-3471
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 3.1.0
<b>Symptom:</b>	Stale BGP Peer-group entry configured under router BGP on SLX Border leaf and Spine devices with none of the BGP neighbors linked with the Peer group.		

<b>Parent Defect ID:</b>	XCO-3471	<b>Issue ID:</b>	XCO-3471
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1. Create a 3-stage CLOS fabric, add devices with MCT leaf, spine, and border-leaf and configure the fabric.</li> <li>2. Convert the 3-stage CLOS fabric to a 5-stage CLOS fabric using the fabric migrate command.  "efa fabric migrate --type "3-to-5-stage" --source-fabric &lt;source-fabric&gt; --destination-3-stage-leaf-spine-pod &lt;pod-name&gt; --destination-3-stage-border-leaf-pod &lt;pod-name&gt;"</li> <li>3. Add super-spine POD devices to the migrated 5-stage CLOS fabric.</li> <li>4. Disconnect the BorderLeaf to Spine links and reconnect the BorderLeaf to Super-Spine links.</li> <li>5. Configure the migrated 5-stage CLOS fabric.</li> </ol>		
<b>Recovery:</b>	Manually delete the stale BGP peer-groups from both the Border Leaf and Spine devices		

<b>Parent Defect ID:</b>	XCO-4128	<b>Issue ID:</b>	XCO-4128
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 3.0.0
<b>Symptom:</b>	Port-channel partial configuration are present on device for SLX NPB devices		
<b>Condition:</b>	Port-channel configuration failed from UI, on device still the partial configuration is present.		
<b>Workaround:</b>	Make sure that all the configuration information are correctly populated from UI so that configuration will not fail on device.		
<b>Recovery:</b>	Login to SLX CLI and delete the given port channel and click on refresh configuration on XCO UI from device action list.		

<b>Parent Defect ID:</b>	XCO-4129	<b>Issue ID:</b>	XCO-4129
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 3.0.0
<b>Symptom:</b>	Disable of vn-tag header strip and enabling of 802.1BR header strip fails from XCO GUI for SLX NPB		
<b>Condition:</b>	When vn-tag header strip is enabled on an interface, disabling the vn-tag header strip and enabling the 802.1BR header strip in a single operation fails from XCO GUI.		

<b>Parent Defect ID:</b>	XCO-4129	<b>Issue ID:</b>	XCO-4129
<b>Workaround:</b>	Disable the vn-tag header strip in first operation (save the port update) and then edit port again for enabling 802.1BR header strip option.		
<b>Recovery:</b>	N/A		

<b>Parent Defect ID:</b>	XCO-6964	<b>Issue ID:</b>	XCO-6964
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.0
<b>Symptom:</b>	Upgrade was successful but a Failed message appears at the end of upgrade.. message is "Status: Failed"		
<b>Condition:</b>	Upgrade XCO with latest build. Issue is seen only on customer setup.		

<b>Parent Defect ID:</b>	XCO-7100	<b>Issue ID:</b>	XCO-7100
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.0
<b>Symptom:</b>	"Target TPVM Version" Does not display until new TPVM is already installed.		
<b>Condition:</b>	During the TPVM upgrade, "Target TPVM Version" gets updated late in the workflow.		
<b>Recovery:</b>	The correct Target TPVM version gets updated after the new version is installed.		

<b>Parent Defect ID:</b>	XCO-8191	<b>Issue ID:</b>	XCO-8191
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	If you run concurrent epg update commands operation as port-group-add or vrf-add on bridge-domain EPGs that are associated with more than one ctag, one or some of the commands may fail with error "Save for device failed".		
<b>Condition:</b>	This is observed more often when more than 3 concurrent EPG port-group-add commands with non-conflicting ports and non-overlapping ctag-range are executed. Occasionally, configuration information that is pushed by one command is not used properly to prepare command recipe for another, causing the failure of one command.		

<b>Parent Defect ID:</b>	XCO-8191	<b>Issue ID:</b>	XCO-8191
<b>Workaround:</b>	Rerunning the failing command will succeed. The error is intermittent and does not cause permanent changes. XCO state information is not affected at any point.		
<b>Recovery:</b>	No recovery is required as no state change is done as part of this failure.		

<b>Parent Defect ID:</b>	XCO-8550	<b>Issue ID:</b>	XCO-8550
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	IPv6 deployment failing with default GW errors even IPv6 XCO server can reach to default Gateway.		
<b>Condition:</b>	IPv6 deployment is failed in below condition. Failed condition- default proto static metric 1024 nexthop via 2600:3c01:e000:e2::2 dev eth0 weight 1 expecting the below pattern - fe80::/64 dev veth0a9acd23 proto kernel metric 256 pref medium default via 2620:100:c:e085:20c:29ff:fee1:3ec1 dev ens160 metric 1024 pref medium		
<b>Workaround:</b>	User has to add the default route using below command- sudo ip -6 route add default via <IPv6 address of default gw> dev <exit interface> example: sudo ip -6 route add default via fc00::5:204:96ff:fed6:f288 dev eth0		

<b>Parent Defect ID:</b>	XCO-8735	<b>Issue ID:</b>	XCO-8735
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	Inventory and device page showing different firmware version.		
<b>Condition:</b>	Post firmware upgrade, Inventory and device page showing different firmware version.		
<b>Workaround:</b>	Post device discovery, device page shows correct firmware ware version.		

<b>Parent Defect ID:</b>	XCO-8829	<b>Issue ID:</b>	XCO-8829
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	New firmware-host registry fails when single quote is used in the password.		

<b>Parent Defect ID:</b>	XCO-8829	<b>Issue ID:</b>	XCO-8829
<b>Condition:</b>	Single quote is used in the password.		
<b>Workaround:</b>	Use the password with no single quote.		

<b>Parent Defect ID:</b>	XCO-9137	<b>Issue ID:</b>	XCO-9137
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.0
<b>Symptom:</b>	EFA upgrade from release 2.7.2 to 3.3.0		
<b>Condition:</b>	When DNS was removed before upgrade.		
<b>Workaround:</b>	DNS configuration should not be changed between upgrades.		
<b>Recovery:</b>	<p>If DNS config removed, after upgrade to EFA 3.3.0 use update_dns.sh script to disallow DNS using following steps:</p> <ol style="list-style-type: none"> <li>1. Bash update-dns.sh --dns-action disallow.</li> <li>2. Get the coredns pod name using k3s kubectl get pods -n kube-system.</li> <li>3. Restart Coredns pod using k3s kubectl delete pod &lt;coredns pod name&gt; -n kube-system.</li> <li>4. Wait for few mins or Restart All efa pods using: <ul style="list-style-type: none"> <li>sudo efactl stop.</li> <li>sudo efactl start.</li> </ul> </li> </ol>		

<b>Parent Defect ID:</b>	XCO-9178	<b>Issue ID:</b>	XCO-9178
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.4.0
<b>Symptom:</b>	When a device is removed from fabric which has QoS configuration, not all QoS configuration is removed from the device and EFA.		
<b>Condition:</b>	Device is deleted from fabric		
<b>Workaround:</b>	User needs to unbind the policies (QoS) from all the relevant targets (fabric/tenant/port/po) before executing the fabric device delete to avoid the stale policies(QoS) in the system.		

<b>Parent Defect ID:</b>	XCO-9190	<b>Issue ID:</b>	XCO-9190
	XCO	<b>Reported in Release:</b>	XCO 3.3.0

Parent Defect ID:	XCO-9190	Issue ID:	XCO-9190
Product:			
Symptom:	VM GUI Library matches shows 2 devices when only 1 device discovered.		
Condition:	Remove all the devices and discover only one device.		

Parent Defect ID:	XCO-9195	Issue ID:	XCO-9195
Product:	XCO	Reported in Release:	XCO 3.3.0
Symptom:	XCO3.3 VM GUI won't allow library copy for matches.		
Condition:	Go to Library page using XCO GUI and try to make a copy using edit option.		

Parent Defect ID:	XCO-9217	Issue ID:	XCO-9217
Product:	XCO	Reported in Release:	XCO 3.3.0
Symptom:	CLI for <b>efact1 restart</b> throwing retrying and failed messages on customer setup only.		
Condition:	Run the CLI <b>efact1 restart</b> command on GTAC setup and user can see the failed pop-up trace even those pods restart completed successfully.		

Parent Defect ID:	XCO-9224	Issue ID:	XCO-9224
Product:	XCO	Reported in Release:	XCO 3.4.0
Symptom:	QoS application is not taking place after changing interface switch port modes or changing the interface from L2 to L3 or vice-a-versa through OOB (out of band) means.		
Condition:	Modifying the switchport mode followed by executing auto DRC.		
Recovery:	Remove all OOB Interface configuration and restore the configuration to the original EFA configured values and re-run DRC.		

Parent Defect ID:	XCO-9270	Issue ID:	XCO-9270
Product:	XCO	Reported in Release:	XCO 3.3.1
Symptom:	Standby node showing 'down' after node reboot		



<b>Parent Defect ID:</b>	XCO-9270	<b>Issue ID:</b>	XCO-9270
<b>Condition:</b>	Trigger SLX reload using below command where the active XCO node is running and check the node status after reboot (when it becomes standby) echo c >/proc/sysrq-trigger		
<b>Recovery:</b>	Reboot both active and standby nodes to recover from the situation		

<b>Parent Defect ID:</b>	XCO-9284	<b>Issue ID:</b>	XCO-9284
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.4.0
<b>Symptom:</b>	Copy default-config to startup-config with maintenance mode enabled will remove all config including QoS policies on a device. Further running DRC does not properly re-install all QoS configuration.		
<b>Recovery:</b>	Remove the device from inventory then register the device back.		

<b>Parent Defect ID:</b>	XCO-9291	<b>Issue ID:</b>	XCO-9291
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.4.0
<b>Symptom:</b>	The fabric internal ports QoS profile is not getting applied on fabric internal ports when Leaf devices are converted from single-homed to multi-homed by adding new Leaf device.		
<b>Condition:</b>	The fabric internal ports QoS profile is not getting applied on fabric internal ports when Leaf devices are converted from single-homed to multi-homed by adding new Leaf device.		
<b>Workaround:</b>	User can issue unbind of fabric internal port QoS profile and rebind the fabric internal port QoS profile using below commands. Unbind Fabric internal ports QoS profile: efa policy qos profile unbind --name <profile_name> --fabric <fabric_name> --port fabric-internal Bind Fabric internal ports QoS profile: efa policy qos profile bind --name <profile_name> --fabric <fabric_name> --port fabric-internal		
<b>Recovery:</b>	User can issue unbind of fabric internal port QoS profile and rebind the fabric internal port QoS profile using below commands. Unbind Fabric internal ports QoS profile: efa policy qos profile unbind --name <profile_name> --fabric <fabric_name> --port fabric-internal Bind Fabric internal ports QoS profile: efa policy qos profile bind --name <profile_name> --fabric <fabric_name> --port fabric-internal		

<b>Parent Defect ID:</b>	XCO-9331	<b>Issue ID:</b>	XCO-9331
	XCO	<b>Reported in Release:</b>	XCO 3.4.0

<b>Parent Defect ID:</b>	XCO-9331	<b>Issue ID:</b>	XCO-9331
<b>Product:</b>			
<b>Symptom:</b>	If a tenant interface level qos profile binding exist on a port channel and the port channel is removed from the device via OOB (Out Of Band) triggering DRC will not re-install the Tenant level interface binding.		
<b>Condition:</b>	Removing a Port Channel from a device via OOB (Out of Band) and triggering DRC.		
<b>Workaround:</b>	Once the port channel is restored by the DRC process on the device the user will need to re-apply/rebind the desired QoS profile on the tenant interface (port channel) using efa policy qos profile bind --name <profile_name> --tenant <tenant_name> --po <port channel ID>		
<b>Recovery:</b>	Once the port channel is restored by the DRC process on the device the user will need to re-apply/rebind the desired QoS profile on the tenant interface (port channel) using efa policy qos profile bind --name <profile_name> --tenant <tenant_name> --po <port channel ID>		

<b>Parent Defect ID:</b>	XCO-9336	<b>Issue ID:</b>	XCO-9336
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.4.0
<b>Symptom:</b>	Inventory device delete is not removing qos config on the spine device		
<b>Condition:</b>	Device deletion from inventory which has QoS configuration		
<b>Workaround:</b>	User needs to unbind the policies (QoS) from all the relevant targets (fabric/tenant/port/po) before running the inventory device delete.		
<b>Recovery:</b>	User needs to unbind the policies (QoS) from all the relevant targets (fabric/tenant/port/po). After this user needs to delete the leftover QoS configuration from SLX.		

<b>Parent Defect ID:</b>	XCO-9341	<b>Issue ID:</b>	XCO-9341
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.2.1
<b>Symptom:</b>	App state for one of the border leaf shows 'cfg-refresh-error'.		

<b>Parent Defect ID:</b>	XCO-9341	<b>Issue ID:</b>	XCO-9341
<b>Condition:</b>	Due to some reason if mariadb restarts on active XCO node, "Error : dial tcp <xco-ip>:3306: connect: connection refused; invalid transaction; invalid transaction" is seen for 'efa fabric show' command. After db connection is successful, the app state for one of the border leaf shows 'cfg-refresh-error'.		
<b>Recovery:</b>	To update the app-state for the device, the below recovery steps can be followed: <ol style="list-style-type: none"> <li>1. On SLX: shut MCT ports.</li> <li>2. On XCO: efa inventory device update -ip &lt;device-ip&gt;.</li> <li>3. On SLX:, no shut MCT ports.</li> <li>4. On XCO: efa inventory device update -ip &lt;device-ip&gt;.</li> </ol>		

<b>Parent Defect ID:</b>	XCO-9354	<b>Issue ID:</b>	XCO-9354
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.3.1
<b>Symptom:</b>	Threshold monitor configuration for monitor types(lif, bfd-session, vxlan-tunnel, mac-table) failed when we configure thru XCO to SLX device(starting from SLX version - 20.5.3).		
<b>Condition:</b>	Use XCO 3.3.1 and SLX version 20.5.3 and try the following cases: <ol style="list-style-type: none"> <li>1. Set monitor threshold for lif, bfd-session, vxlan-tunnel, mac-table thru XCO failed only when we use count/interval fields.</li> <li>2. Unset monitor threshold for lif, bfd-session, vxlan-tunnel, mac-table thru XCO failed. Here Unset logic tries to clear all fields including count/interval fields by default.</li> </ol>		
<b>Workaround:</b>	<ol style="list-style-type: none"> <li>1. Set monitor threshold for types lif, bfd-session, vxlan-tunnel, &amp; mac-table thru XCO without count/interval fields.</li> <li>2. Unset can't be done by XCO. Instead it can be unset directly in SLX using the following commands:   no threshold-monitor &lt;monitor-type&gt;   Or use XCO 3.3.1 with SLX version 20.5.2a for functioning without any issue.</li> </ol>		

<b>Parent Defect ID:</b>	XCO-9362	<b>Issue ID:</b>	XCO-9362
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.4.0
<b>Symptom:</b>	The fabric internal ports QoS profile is not getting applied on intended ports: <ol style="list-style-type: none"> <li>1) When a new device is being added to CLOS fabric and fabric is configured.</li> <li>2) When a new rack is added to non-CLOS fabric and fabric is configured.</li> </ol>		

<b>Parent Defect ID:</b>	XCO-9362	<b>Issue ID:</b>	XCO-9362
<b>Condition:</b>	<p>Pre-condition: Fabric internal ports QoS profile is already applied on a fabric (CLOS or non-CLOS).</p> <p>Issue will be seen:</p> <ol style="list-style-type: none"> <li>1) When a new device is being added to CLOS fabric and fabric is configured.</li> <li>2) When a new rack is added to non-CLOS fabric and fabric is configured</li> </ol>		
<b>Workaround:</b>	<p>User can issue unbind of fabric internal port QoS profile and rebind the fabric internal port QoS profile using below commands.</p> <p>Unbind Fabric internal QoS profile: efa policy qos profile unbind --name &lt;profile_name&gt; --fabric &lt;fabric_name&gt; --port fabric-internal</p> <p>Bind Fabric internal QoS profile: efa policy qos profile bind --name &lt;profile_name&gt; --fabric &lt;fabric_name&gt; --port fabric-internal</p>		
<b>Recovery:</b>	<p>User can issue unbind of fabric internal port QoS profile and rebind the fabric internal port QoS profile using below commands.</p> <p>Unbind Fabric internal QoS profile: efa policy qos profile unbind --name &lt;profile_name&gt; --fabric &lt;fabric_name&gt; --port fabric-internal</p> <p>Bind Fabric internal QoS profile: efa policy qos profile bind --name &lt;profile_name&gt; --fabric &lt;fabric_name&gt; --port fabric-internal</p>		

<b>Parent Defect ID:</b>	XCO-9363	<b>Issue ID:</b>	XCO-9363
<b>Product:</b>	XCO	<b>Reported in Release:</b>	XCO 3.4.0
<b>Symptom:</b>	Tenant and Tenant Interface QoS policy information is not fully removed from EFA DB and SLX device when the tenant delete with --force option is specified.		
<b>Condition:</b>	Issue is observed when user issues the command 'efa tenant delete --name <tenant_name> --force'		
<b>Workaround:</b>	This is covered in the targeted Generic Release notes (item) User needs to unbind the policies (QoS) from all the relevant targets (fabric/tenant/port/port channel/tenant endpoint group) before executing the force operations including delete to avoid the stale policies(QoS) in the system.		

<b>Parent Defect ID:</b>	XCO-9381	<b>Issue ID:</b>	XCO-9381
<b>Product:</b>	XCO	<b>Reported in Release:</b>	EFA 2.7.2
<b>Symptom:</b>	9740 devices with breakout port configured, DRC fails for even numbered port.		

Parent Defect ID:	XCO-9381	Issue ID:	XCO-9381
Condition:	If XCO is upgraded from previous version to 3.2.0 version.		
Workaround:	Perform fresh install followed by reconfiguration of breakout ports and its respective configuration.		

## 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](http://www.extremenetworks.com/support/contact).

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

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

## Subscribe to Product Announcements

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

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.

4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

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