



ExtremeCloud™ Orchestrator v4.0.2 Troubleshooting Reference Guide

Comprehensive Solutions and Management

9039316-01 Rev AA
April 2026



Copyright © 2026 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: <https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks>

Open Source Declarations

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

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




Table of Contents

Abstract.....	vii
Preface.....	viii
Text Conventions.....	viii
Documentation and Training.....	ix
Open Source Declarations.....	x
Training.....	x
Help and Support.....	x
Subscribe to Product Announcements.....	xi
Send Feedback.....	xi
What's New in this Document.....	12
Troubleshooting Backup and Restore.....	13
EFA or XCO Backup Authentication Fails.....	13
Interface Config Missing After Restore from EFA or XCO Backup.....	13
Interface Set-Speed Config Missing in Running Config.....	15
System Restore Operation Failures.....	15
Error: Restore Operation Has Failed, Exit Status 1.....	15
Error: Restore Operation Has Failed, Exit Status 255.....	16
Supportsave Fails to Collect Data.....	16
Unable to Run EFA or XCO System Backup.....	17
Troubleshooting Certificate Management.....	18
Applying Third-party Certificate Acquired Through Trusted CAs to EFA or XCO Throws Error.....	18
Certificate Mismatch Error on the SLX Console.....	18
Certificate Troubleshooting.....	19
SLX Device Shows Expired Crypto CA Certificates.....	20
Troubleshooting Device Management.....	22
Configuration Restore Functionality is Not Working as Expected.....	22
Drift and Reconcile Shows 'In-Progress' Status Despite 'Failed' Status for Several Devices.....	24
Drift and Reconcile has Not Started, is Incomplete, or SLX is Stuck in Maintenance Mode Following DRC.....	25
Drift and Reconcile Resets the Timezone on the SLX.....	25
Debug Device Settings Changes.....	26
EFA Firmware-Download Prepare Add Command Fails.....	27
EFA Inventory Config-Replay Execute Command Fails.....	27
Execution Error: 503 Service Unavailable.....	28
HTTPS Connection Issues Between XCO and SLX.....	28
HTTP/HTTPS Process Error and Multiple Defunct Process.....	28
Leaf and Spine Nodes Stuck in "cfg-refresh error" After Restore.....	30
SNMPv3 User Credentials Issue with Special Character (\$) in Password.....	31

Switch Replacement Procedure Fails.....	33
Verify XCO is Registered to Receive Traps from SLX.....	33
Verify the Types of Device Updates Done.....	33
Minimal Update Message.....	34
Full Update Message.....	34
Verify the Device Supports Minimal Updates.....	34
Verify Device Update Notification on Out of Band Config Change on SLX.....	35
Verify Out of Band Config Change (Non-XCO) Triggered XCO for a Full Update.....	35
Verify Reachability of XCO to a New Management IP Address.....	35
Verify Switch Registration Process Completion on XCO Using REST APIs.....	35
Troubleshooting Fabric Skills.....	37
Access Token Expires Before the Specified Expiration Time.....	38
Active XCO or EFA Node is Down.....	39
Cannot Ping IP Address of Sub-Interface.....	40
Check Status of Restore API.....	40
Daemonset "goraslog-service" is Stuck in a "Not Ready" State.....	40
Devices in CFG Refreshed or CFG Refreshed Error State.....	41
EFA or XCO Device Status is "cfg refreshed" or "cfg refresh error".....	41
EFA Still Reporting "cfg refreshed" Post Drift Reconcile.....	42
"Error: dial tcp" When Running "efa fabric show" Command.....	44
Device in "cfg refresh error" After MCT CCP Restoration.....	45
Multiple or Duplicate IPs on Fabric Interfaces.....	45
App-state Showing 'cfg-refreshed' for OOB L2VPN.....	46
Device Showing "cfg refreshed" After Annual Interface Description Update.....	46
Device in "cfg refresh error" After BMC Upgrade	47
EFA Fabric Show Command Output Showing Nodes with "cfg refreshed" App State.....	47
XCO Fabric Device Displays "cfg refreshed" After Manual Update of Interface Description.....	48
Device Addition to Fabric Fails.....	49
Switch Addition to Fabric Fails.....	49
Addition of Multiple Devices to Fabric Fails.....	49
EFA Status Shows Nothing and gofaultmanager and gopolicy POD Services Stuck in init State.....	50
EFA Commands Fail to Execute.....	51
EFA Commands Fails with "Error 408".....	52
EFA Command Fails with "Dial TCP xx.xx.xx.xx:80: Connect: Connection Refused" Message.....	52
EFA Error When Registering a RELP Handler on Endpoint Using FQDN.....	52
EFA Fabric Show Command Fails After Active Node Restart Causing CFG Refreshed Error on Devices.....	53
EFA or XCO is Down on Standby TPVM	54
EFA Fabric health Not in Sync with EFA Fabric Topology Show.....	56
EFA or XCO VM is down and Came Up After Stopping the Services.....	57
EFA or XCO Execute-CLI Reports SLX Switches as Unreachable.....	57
EFA or XCO Execute-CLI Fails on SLX.....	58
EFA or XCO Fails to Resolve Non-Cluster FQDN.....	58
EFA or XCO Fails to Log in to Switch.....	59
Fabric Skill Troubleshooting.....	60


Fabric Skill is Down After a Failed Token Certificate Renewal.....	62
Fabric Configuration Fails with EFA or XCO.....	65
How to Add a Second Link to an Existing Fabric via EFA or XCO.....	66
Log in to EFA or XCO Fails.....	66
LDAP Authentication Configuration Using FQDN Fails.....	67
Management Subinterface IP Address Not Listed.....	67
Recover a Leaf Node after Deletion.....	68
SLX, EFA, or XCO Config Missing After SLX Reload.....	69
TPVM Unable to Authenticate to Red Hat Directory Service (LDAP).....	70
Unable to Run EFA Commands.....	70
Unable to Connect to the Server.....	71
Update Authentication Preference.....	72
Update of Maximum Password Age of SLX Password Fails.....	73
Analyze 10 Second Traffic Loss During DN Leaf Power Cycle.....	73
XCO or EFA Authentication Stops Working when Users Configure a Different Authentication.....	74
EFA CLI Login Fails with Error "Failed to Fire Hook".....	76
Troubleshooting Installation and Deployment.....	77
Connection to Server Localhost: 8080 Fails.....	77
EFA or XCO OVA Image Reports Read Only File System.....	78
EFA or XCO Deployment Fails on TPVM.....	78
EFA or XCO Does Not Start Properly and the PODs Stuck in Init State.....	79
Identify the Active Node that Serves as the Database for Kubernetes Clusters.....	80
Installation of Database (MariaDB) Fails.....	80
Installation of XCO or EFA Fails.....	80
Login to XCO CLI Fails.....	80
Pod is in a Crashloopback State.....	81
Post-Uninstallation of EFA or XCO, Existing Fabric Configuration Persists on SLX.....	81
RabbitMQ Pods Continuously being Deleted and Redeployed or in CrashLoopBackOff....	82
Split Brain Issue in EFA or XCO HA Deployment.....	83
TPVM Deployment Fails.....	84
Transport Endpoint is Not Connected.....	85
Uninstall a Failed or Partial Installation.....	86
XCO Deployment with Management Sub-Interface Fails.....	86
XCO Instability After Configuring TPVM LDAP Host.....	87
Troubleshooting License Management.....	89
License is Not Properly Installed.....	89
License Error Handling.....	90
License File Already Exists.....	90
License Expired.....	90
License Parsing Error.....	90
Not an XCO License.....	91
Troubleshooting Network Infrastructure Components.....	92
RabbitMQ Log Rotation Does Not Remove or Archive Old Logs.....	92
Two Instances of Each Service is Seen in EFA or XCO.....	92
Verify Reachability of XCO to a New Management IP Address.....	93
Troubleshooting Policy Services.....	94
QoS Profiles and Maps are not Automatically Removed.....	94

Troubleshooting Tenant Services.....	95
Adding Multiple BFD and Static Routes to SLX Config via XCO Causes a Panic on SLX.....	95
APS Admin Down or Admin Up Operations Fail.....	96
APS Admin Down Operation Fails.....	96
APS Admin Up Operation Fails.....	96
BGP Peer Group Creation Fails Due to Permanent EFA Error.....	96
Change Anycast MAC Address in SLX Fabric Without Impacting Services.....	97
EPG Creation or Update Fails.....	97
EFA or XCO Tenant Command Fails.....	98
Failed to Save Devices During EPG Creation.....	98
How to Enable BFD Under a VE Interface via XCO.....	99
MCT Cluster Config Missing Between Rack Devices Error Message is Not Relayed in REST API.....	100
Static Routes Persist on SLX Devices after Deletion from XCO.....	101
Successful Tenant Creation with Overlapping L2 and L3 VNIs.....	101
Tenant Creation Fails.....	102
Tenant Creation Fails with Port Error	102
Tenant Creation Fails Intermittently After Initial Fabric Creation: Ports Unavailable.....	102
Tenant Service Restart Scenario.....	103
Scenario 1: Missing RabbitMQ Readiness Check.....	103
Scenario 2: Liveness Probe Failure During Long Initialization.....	104
Tenant VRF Update Fails.....	105
Unable to Delete an EPG with Comma in Name.....	106
VLAN Configuration Fails When Adding Port or PO to EPG.....	107
VRF Already Exists on Device Error	108
Troubleshooting Upgrade and Migration.....	111
Active Node is Down after EFA or XCO Upgrade.....	111
EFA or XCO Upgrade Fails on Ubuntu Server.....	112
EFA or XCO Upgrade to 3.2.x Fails.....	112
Incorrect EFA or XCO Version Displayed After Upgrade.....	113
Kubernetes POD CoreDNS Stuck in CrashLoopBackOff Post Upgrade	113
Login to EFA or XCO Fails after Upgrade.....	114
SLX Upgrade Failure and Inability to Restore.....	116
TPVM Incremental Upgrade Fails.....	117
TPVM Upgrade on the Standby Node Fails During a Data Fabric Upgrade.....	117
User Accounts Deleted Post TPVM Upgrade.....	118
XCO Firmware Upgrade Fails.....	118
XCO or EFA Lost Connection to Devices After Upgrade.....	118
Scenario 1: Connection Loss and Certificate Error	119
Scenario 2: Fallback Communication Failure.....	119
EFA Status Remains Down After Upgrade.....	120



Abstract

This troubleshooting reference guide for ExtremeCloud™ Orchestrator version 4.0.2 provides comprehensive technical instructions for diagnosing and resolving issues across backup and restore operations, certificate and device management, fabric skills, installation and deployment, licensing, and upgrade scenarios. Designed for IT professionals managing SLX devices and XCO environments, the guide addresses problems such as authentication failures during system backup, configuration drift and reconcile inconsistencies, certificate mismatches, and pod initialization errors. It includes structured procedures for resolving HTTPS connectivity issues, SNMPv3 credential conflicts, and fabric configuration anomalies, with emphasis on command-line diagnostics, REST API usage, and system log analysis. The guide also outlines upgrade and migration challenges, including token expiry misconfigurations, TPVM deployment failures, and post-upgrade service disruptions. Troubleshooting scenarios are supported with examples, error messages, and corrective actions, ensuring technical readers can maintain operational integrity and configuration consistency across ExtremeCloud deployments.



Preface

Read the following topics to learn about:

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

Text Conventions

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

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

Table 1: Notes and warnings






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

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
<i>Words in italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic text</i>	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member [member...]</i> .
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and Software Compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other Resources](#) such as articles, white papers, and case studies

Open Source Declarations

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

Training

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

Help and Support

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

Extreme Portal

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

The Hub

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

Call GTAC

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

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

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

Subscribe to Product Announcements

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

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

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


Send Feedback

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

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

To send feedback, email us at Product-Documentation@extremenetworks.com.

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



What's New in this Document

The following table describes information added to this guide for the ExtremeCloud Orchestrator (XCO) 4.0.2 software release. For more information about this release, see ExtremeCloud Orchestrator Release Notes, 4.0.2.

Table 4: Summary of changes

Description	Link
A new topic, "Tenant Service Restart Scenario", explains why Tenant Services might restart repeatedly during system restore, along with steps to resolve the issue.	Tenant Service Restart Scenario on page 103
A new topic "Interface Set-Speed Config Missing in Running Config" describes troubleshoot when the interface Set-Speed config is missing in running config.	Interface Set-Speed Config Missing in Running Config on page 15
A new topic "Error: Restore Operation Has Failed, Exit Status 255" describes a system restore failure case.	System Restore Operation Failures on page 15
New topic "EFA CLI Login Fails with Error "Failed to Fire Hook"" describes troubleshooting steps when the CLI login fails with the error "Failed to Fire Hook"	EFA CLI Login Fails with Error "Failed to Fire Hook" on page 76



Note

From XCO 3.2.0 onwards, 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.



Troubleshooting Backup and Restore

[EFA or XCO Backup Authentication Fails](#) on page 13

[Interface Config Missing After Restore from EFA or XCO Backup](#) on page 13

[Interface Set-Speed Config Missing in Running Config](#) on page 15

[System Restore Operation Failures](#) on page 15

[Supportsave Fails to Collect Data](#) on page 16

[Unable to Run EFA or XCO System Backup](#) on page 17

Learn about troubleshooting backup and restore issues.

EFA or XCO Backup Authentication Fails

An authentication failure occurred during the EFA or XCO backup process, causing the backup to fail.

```
(efa:efabackup)efabackup@KVDCxxxx-cTPVM02:~$ efa system backup --fabric-name XXXYYY-  
Fabric --remote
```

```
Generating backup of EFA...
```

Backup Errors:

```
ssh: handshake failed: ssh: unable to authenticate, attempted methods [none  
password], no supported methods remain Backup Location: /apps/efa_logs/backup/  
EFA-2023-07-27T15-59-04.706.tar
```

```
{"@time":"2023-07-26T16:01:12.091424  
+08","App":"efa","level":"error","msg":"Failed","reason":"invalid credentials","request":  
{ "cmd":"auth:create-client-access-token","params":{"ClientID":"3377a122-37ad-45ac-9e73-  
a84f9553402b","UserName":"efabackup"}}, "rqId":"cca0219d-7e79-45b1-8407-a258f800233e"}
```

The issue occurred because an authentication token has expired.

Resolution

Log out and log back into EFA or XCO.



Note

You can set the token duration to a minimum of 1 hour and a maximum of 8 hours, with the default setting of 1 hour.

Interface Config Missing After Restore from EFA or XCO Backup

After restore from a previous backup to resolve an issue, the interface (for example, 0/12) config was missing on the switch (for example, switch2).

The following configurations found on Switch1:

```
interface Ethernet 0/12:1
  no shutdown
!
interface Ethernet 0/12:2
  no shutdown
!
interface Ethernet 0/12:3
  no shutdown
!
interface Ethernet 0/12:4
  no shutdown
```

The following configurations found on Switch2:

```
interface Ethernet 0/12
  no shutdown
```

The following are the configurations on EFA or XCO:

```
efa tenant create --name "cnis_upf_tenant" --description "cnis_upf_tenant" --type private
--vlan-range 801,805,3220,3608,3609,201-210,251-260,820-860,220-250,270-300,811-819 --vrf-
count 12 --port
192.168.246.22[0/1-4,0/6-7,0/10:1-4,0/11:1-4,0/13:1-4,0/5:1-4,0/8:1-4,0/9:1-4],192.168.246
.21[0/1-4,0/6-7,0/10:1-4,0/11:1-4,0/12:1-4,0/13:1-4,0/5:1-4,0/8:1-4,0/9:1-4]
```

The issue is related to the EFA or XCO configuration, where some interfaces are not managed (cfg-not-managed) while others are in sync (cfg-in-sync). Upon running the efa inventory device interface list-breakout --ip=10.64.208.28 command to check the interface breakout configuration, it was found that certain interfaces (for example, 0/12:1 to 0/12:4) were marked as cfg-in-sync, indicating they were configured via XCO. In contrast, some interfaces were marked as cfg-not-managed, meaning they were not managed by XCO and were likely imported from the SLX configuration.

```
(efa:extreme)extreme@tpvm-112:~$ efa inventory device interface list-breakout --ip=10.64.208.28
+-----+-----+-----+
| IP Address | Name | AppState |
+-----+-----+-----+
| 10.64.208.28 | 0/10:1 | cfg-not-managed | ==> oob [imported from slx config hence cfg-not-managed]
+-----+-----+-----+
| | 0/10:2 | cfg-not-managed |
+-----+-----+-----+
| | 0/10:3 | cfg-not-managed |
+-----+-----+-----+
| | 0/10:4 | cfg-not-managed |
+-----+-----+-----+
| | 0/12:1 | cfg-in-sync | ==> non oob [configured through xco cfg-in-sync]
+-----+-----+-----+
| | 0/12:2 | cfg-in-sync |
+-----+-----+-----+
| | 0/12:3 | cfg-in-sync |
+-----+-----+-----+
| | 0/12:4 | cfg-in-sync |
+-----+-----+-----+
| | 0/4:1 | cfg-in-sync |
+-----+-----+-----+
| | 0/4:2 | cfg-in-sync |
+-----+-----+-----+
| | 0/4:3 | cfg-in-sync |
+-----+-----+-----+
| | 0/4:4 | cfg-in-sync |
```

```

+ +-----+-----+
| | 0/8:1 | cfg-not-managed |
+ +-----+-----+
| | 0/8:2 | cfg-not-managed |
+ +-----+-----+
| | 0/8:3 | cfg-not-managed |
+ +-----+-----+
| | 0/8:4 | cfg-not-managed |
+ +-----+-----+
Interface Details
--- Time Elapsed: 104.013201ms ---

```

Resolution

Manually remove the configurations from SLX to permanently delete the breakout configurations.



Note

- Non-OOB configurations previously set via XCO on SLX will be restored after DRC with reconciliation, as XCO serves as the authoritative source.
- If OOB breakout configurations are deleted, they are removed from the breakout table in XCO after an inventory update, since they were not configured through XCO. This happens because they were initially imported from SLX, and their removal from SLX triggers their deletion in XCO.
- As a best practice, use the XCO command **efa inventory device interface set-breakout** to configure port breakout settings on the SLX.

Interface Set-Speed Config Missing in Running Config

Missing set-speed config in running config causes backup restore failure. This happens when using unqualified optics and breakout configs.

Breakout needs qualified optics (for example, 100G optics → 4x25G only, 40G optics → 4x10G only). For unqualified optics, ensure to manually set port speed before DB restore, Restoring from default config without setting port speed is not supported.

System Restore Operation Failures

Use this topic to learn about troubleshooting various restore failure scenarios.

Error: Restore Operation Has Failed, Exit Status 1

The EFA or XCO system restore backup fails with the error "Restore operation has failed, exit status 1". This issue occurs when attempting to restore a backup taken from a dual-server setup (active and standby) on a single server (standby) with a different IP stack configuration.

```

efa:slx@WEXTREME:~/XCO_images/efa$ efa system restore --backup-tar EFA-3.3.0-
GA-2023-10-01T00-00-15.222.tar
Performing EFA restore using EFA-3.3.0-GA-2023-10-01T00-00-15.222.tar

```


IP Address	Status	Reason
21.151.151.102	Failed	Supportsave process didnt start on the device 21.151.151.102
21.151.151.101	Failed	Supportsave process didnt start on the device 21.151.151.101

The issue is caused by the permission settings on the external server.

Warning: SLX9250-32C, Copy support upload operation failed. Reason: Please check the username, password, or directory/file permission.

Resolution

Correct the permission settings on the external server.

Unable to Run EFA or XCO System Backup


The attempt to run the **efa system backup** command returns the following error message:

```
Failed to execute service lock API due to error: Service Approval Failed for Process BACKUP.
```

As the message queue for the fabric service continues to grow, issues may eventually arise with the system backup.

Resolution

Try purging the message bus to resolve the blockage. Contact GTAC for help with checking and purging the RabbitMQ queues.



Troubleshooting Certificate Management

[Applying Third-party Certificate Acquired Through Trusted CAs to EFA or XCO Throws Error](#) on page 18

[Certificate Mismatch Error on the SLX Console](#) on page 18

[Certificate Troubleshooting](#) on page 19

[SLX Device Shows Expired Crypto CA Certificates](#) on page 20

Learn about troubleshooting certificate management issues.

Applying Third-party Certificate Acquired Through Trusted CAs to EFA or XCO Throws Error

When applying a third-party certificate acquired through trusted Certificate Authorities (CAs) to the EFA or XCO, the following error occurs:

```
(efa:administrator)administrator@ubuntu-svr1804-ta1:~/efa$ efa certificate server --
certificate ca.pem --key ca-key.pem
Please wait as the certificates are being installed...
[sudo] password for administrator:
Error message during certificate installation: Using /etc/efa/efa.conf for EFA settings
Error updating traefik with efasecret
```

The current EFA or XCO implementation requires the Subject Alternative Name (SAN) to be "efa.extremenetworks.com." If any other SAN is used, the process will result in an error.

Resolution

Ensure that the server certificate is generated with "efa.extremenetworks.com" as the Subject Alternative Name.

Certificate Mismatch Error on the SLX Console

Following a fresh installation of XCO and restoration of a backup, certificate mismatch errors occur on the SLX console, preventing XCO from logging into the SLX.

```
2023/10/06-19:01:42, [SEC-3112], 282017,, INFO, SLX, Event: X509v3, Certificate
Validation failed, Info: Reason = unable to get local issuer certificate,
Certificate Details = [Subject CN efa.extremenetworks.com, Serial
206813760996099062386055851929706211376239302260 Issuer /C=US/ST=CA/O=Extreme Networks/
```

```
OU=Extreme Fabric Automation Intermediate/CN=EFA Intermediate CA/
emailAddress=support@extremenetworks.
```

The issue is due to a new fingerprint generated during the fresh install, requiring re-installation of all certificates on the SLX.

Resolution

To resolve the certificate mismatch issue on the SLX console, complete the following steps after restoring the backup on XCO:

1. Run the **efa certificate device install --ip <ip_address> --force** command on XCO to update HTTPS and OAUTH certificates on the SLX.
2. Run the **efa inventory device update --ip <ip_address>** command on XCO to update the syslog certificate on the SLX.

Certificate Troubleshooting

Issue	Resolution
<p>My device is registered but the certificates do not appear on the SLX device.</p>	<p>Try the following:</p> <ul style="list-style-type: none"> • Ensure that the device is running at least SLX-OS 20.1.x. • Ensure that the time on the SLX device and the time on the Extreme Fabric Automation host device are synchronized. • Ensure that the certificates are installed. Run the efa certificate device install command.
<p>How do I know about the certificate expiry in XCO?</p>	<ul style="list-style-type: none"> • Run the following REST API to get the expiry date of all the certificates of XCO: <pre>curl -X GET 'https://<vip>:8078/v1/monitor/certificate/expiry' --header 'Authorization:Bearer eyJhbGciOiJSUzI... '</pre> • Run the following openssl command: <pre>extreme@tpvm:~\$ openssl x509 -in <Location of the certificate> -noout -enddate</pre> • Run the efa certificate expiry show command.

Issue	Resolution
<p>How do I verify the certificate provided by XCO through its ingress interface?</p>	<p>Run the following command. The output should indicate that <code>efa.extremenetworks.com</code> is present.</p> <pre>\$ openssl s_client -connect <EFA_IP_ADDR>:443</pre>
<p>There is a security violation on the switch when XCO (installed on TPVM) logs in and tries to access the switch with different usernames. You observe the following logs on SLX console:</p> <pre>1018 AUDIT, 2021/10/14-17:26:57 (GMT), [SEC-3021], INFO, SECURITY, extreme/root/ 10.20.32.141/ssh/CLI,, SLX, Event: login, Status: failed, Info: Failed login attempt through REMOTE, IP Addr: 10.20.32.141 1017 AUDIT, 2021/10/14-17:26:55 (GMT), [SEC-3020], INFO, SECURITY, admin/admin/ 10.20.32.141/ssh/CLI,, SLX8720-32C, Event: login, Status: success, Info: Successful login attempt via REMOTE, IP Addr: 10.20.32.141 1002 AUDIT, 2021/10/14-17:26:41 (GMT), [SEC-3020], INFO, SECURITY, admin/admin/ 10.20.32.141/ssh/CLI,, SLX8720-32C, Event: login, Status: success, Info: Successful login attempt via NETCONF, IP Addr: 10.20.32.141</pre>	<p>Try the following:</p> <ul style="list-style-type: none"> • Ensure that you have correctly followed the system restore process. • Ensure that all the devices are registered. • Ensure that the certificates are installed on the devices to enable secure connections. Run the <code>efa certificate device install --ips <ip-addr> certType [http token]</code> command to install the HTTPS or OAuth2 certificate on one or more devices.

SLX Device Shows Expired Crypto CA Certificates

The SLX device displays expired Crypto CA certificates, and updating them from XCO fails. This prevents configuration changes from being made on SLX from XCO.

The expired certificates on SLX cause communication failures between SLX and XCO.

Resolution

1. Manually delete the expired certificates on SLX.

```
no crypto ca import-pkcs type pkcs12 cert-type https
```

2. Disable HTTP server.

```
http server use-vrf mgmt-vrf shutdown
```

3. Verify the HTTP server status.

```
show http server status
VRF-Name: default-vrf Status: HTTP Disabled and HTTPS Disabled
VRF-Name: mgmt-vrf Status: HTTP Enabled and HTTPS Disabled
```

4. Update the SLX username and password from XCO.

```
efa inventory device update --ip xx.xxx.xx.xx --username admin --password xxxxxx
```

5. Install the new certificate.

```
efa certificate device install --ip 192.xx.xx.xx --cert-type https
```



Note

- Ensure HTTP server is enabled on both management and default VRF after installing the certificate.

```
show http server status
VRF-Name: default-vrf Status: HTTP Enabled and HTTPS Enabled
VRF-Name: mgmt-vrf Status: HTTP Disabled and HTTPS Enabled
```


- If HTTP is disabled, restart it on the affected (disabled) VRF.

```
show http server status
VRF-Name: default-vrf Status: HTTP Enabled and HTTPS Disabled
VRF-Name: mgmt-vrf Status: HTTP Disabled and HTTPS Enabled
```

```
PD2_L01(config)#http server use-vrf default-vrf shutdown
PD2_L01(config)#no http server use-vrf default-vrf shutdown
```

- If the device status shows `cfg`, run a drift reconciliation (DRC) command.

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



Troubleshooting Device Management

- [Configuration Restore Functionality is Not Working as Expected on page 22](#)
- [Drift and Reconcile Shows 'In-Progress' Status Despite 'Failed' Status for Several Devices on page 24](#)
- [Drift and Reconcile has Not Started, is Incomplete, or SLX is Stuck in Maintenance Mode Following DRC on page 25](#)
- [Drift and Reconcile Resets the Timezone on the SLX on page 25](#)
- [Debug Device Settings Changes on page 26](#)
- [EFA Firmware-Download Prepare Add Command Fails on page 27](#)
- [EFA Inventory Config-Replay Execute Command Fails on page 27](#)
- [Execution Error: 503 Service Unavailable on page 28](#)
- [HTTPS Connection Issues Between XCO and SLX on page 28](#)
- [Leaf and Spine Nodes Stuck in "cfg-refresh error" After Restore on page 30](#)
- [SNMPv3 User Credentials Issue with Special Character \(\\$\) in Password on page 31](#)
- [Switch Replacement Procedure Fails on page 33](#)
- [Verify XCO is Registered to Receive Traps from SLX on page 33](#)
- [Verify the Types of Device Updates Done on page 33](#)
- [Verify the Device Supports Minimal Updates on page 34](#)
- [Verify Device Update Notification on Out of Band Config Change on SLX on page 35](#)
- [Verify Out of Band Config Change \(Non-XCO\) Triggered XCO for a Full Update on page 35](#)
- [Verify Reachability of XCO to a New Management IP Address on page 35](#)
- [Verify Switch Registration Process Completion on XCO Using REST APIs on page 35](#)

Learn about troubleshooting device management issues.

Configuration Restore Functionality is Not Working as Expected

The automated test failed because the config-replay (**`efa inventory config-replay detail --uuid <uuid-number>`**) command took longer than expected to complete, resulting in loss of connectivity between EFA or XCO and SLX. It remained in progress beyond the test's sleep time, which was set for 7 minutes.

The following steps outline the scripted event sequence:

1. Run the **config-replay** command on a device.

```
efa inventory config-replay execute --ip=x.x.x.x --startup-config
```

2. Run the following command to verify the config-replay status.

```
efa inventory config-replay detail --uuid <uuid-number>
```

The following are the log details:

```
24-03-05 14:53:19.160 Test Step 20: Restore configuration on <host> now hosting standby
EFA - 0:7:17.021
24-03-05 14:53:19.161 Sub Test Step 20_1: Execute config-replay and check details -
0:7:17.008
24-03-05 14:53:23.651 Executed command: 'efa inventory config-replay execute --ip
21.151.151.102 --uuid 5b9f6d7c-d841-4669-be35-782c331e18c5 --startup-config'
24-03-05 14:53:25.097 Sleeping for 120 seconds...
24-03-05 14:55:25.102 Executed command: 'efa inventory config-replay detail --uuid
4dbe6a6b-aff9-43ea-9e59-9b01643c6964'
...
24-03-05 15:00:36.173 Assertion Failed: Config-replay was not successful
```

Despite the failure during the automated test, a manual check later showed that the command completed successfully in about five minutes.

```
(efa:extreme)extreme@host-name:~$ efa inventory config-replay detail --uuid 4dbe6a6b-
aff9-43ea-9e59-9b01643c6964
-----+-----+
| NAME                | VALUE                |
-----+-----+
| UUID                | 4dbe6a6b-aff9-43ea-9e59-9b01643c6964 |
-----+-----+
| Device IP          | x.x.x.x              |
-----+-----+
| Status           | success           |
-----+-----+
| Config Backup UUID | 5b9f6d7c-d841-4669-be35-782c331e18c5 |
-----+-----+
| Config Backup SSID | 1                    |
-----+-----+
| Execution Reason   | manual               |
-----+-----+
| operation          | manual trigger       |
-----+-----+
| Start Time         | 2024-03-05 14:53:24 +0100 CET |
-----+-----+
| Last Modified Time | 2024-03-05 14:58:10 +0100 CET |
-----+-----+
| Duration        | 4m45.757834235s  |
-----+-----+
--- Time Elapsed: 100.021368ms ---
(efa:extreme)extreme@host-name:~$
```

The **config-replay** command was run with the `-startup-config` flag, requiring a device reload. This caused a loss of connectivity between EFA or XCO and SLX during the script's wait time.

Resolution

- Extend the automated script polling interval for the config-replay status.

- Suspend the test case until device connectivity is restored, then proceed. This can be achieved by monitoring the device reload status rather than the config-replay status in the test suite.

Drift and Reconcile Shows 'In-Progress' Status Despite 'Failed' Status for Several Devices

Drift and Reconcile (DRC) status remains 'in-progress' even after 4 days, despite showing 'failed' for some devices.

The issue occurs due to device discovery failure during inventory reconcile.

```
(efa:admin)]$ efa inventory drift-reconcile detail --uuid=e7f794a8-c167-4e5a-b4c6-0f89aeaa3233
+-----+
| NAME                | VALUE                |
+-----+
| UUID                | e7f794a8-c167-4e5a-b4c6-0f89aeaa3233 |
+-----+
| Device IP          | 1.2.3.4              |
+-----+
| Status              | in-progress          |
+-----+
| Failure Reason    | device discovery failed After |
|                    | Inventory Reconcile         |
+-----+
| Execution Reason    | manual               |
+-----+
| operation            | drift-and-reconcile |
+-----+
| Inventory Status    | inventory-dr-success |
+-----+
| Is Inventory config Refreshed | true                |
+-----+
| Inventory Duration  | 25.408171436s       |
+-----+
| Fabric Status       | unknown              |
+-----+
| Is Fabric config Refreshed | false              |
+-----+
| Fabric Duration     |                      |
+-----+
| Policy Status       | unknown              |
+-----+
| Is Policy config Refreshed | false              |
+-----+
| Policy Duration     |                      |
+-----+
| Tenant Status       | unknown              |
+-----+
| Is Tenant config Refreshed | false              |
+-----+
| Tenant Duration     |                      |
+-----+
| Device Update Count | 2                    |
+-----+
| Device Update Total Duration | 3m3.675810243s    |
+-----+
| Maintenance Mode Disable |                      |
| Duration            |                      |
+-----+
```

Start Time	2024-02-12 12:45:55 +0100 CET	
+-----+		
Last Modified	2024-02-12 12:49:37 +0100 CET	
+-----+		
Duration	3m42.308698216s	
+-----+		

A certificate error is found in the logs:

```
{"@time":"2024-02-15T11:42:49.54647
CET", "App":"inventory", "Device":"1.2.3.4", "UseCase":"Trigger Deep device
update", "level":"error", "msg":"AddSyslogServerEntry: Failed to compare local syslog SHA1
with the syslog certs on
```

Resolution

Re-install the certificates using the following command:

```
efa certificate device install --ip 1.2.3.4
```

Drift and Reconcile has Not Started, is Incomplete, or SLX is Stuck in Maintenance Mode Following DRC

Verify the following:

1. Ensure that the goraslog service is running in XCO.
2. Check the file at the <path>/raslog/raslog-server.log location to see if XCO has received the SMAN/DCM-1116 messages from SLX.
3. Check the DRC state in XCO using the **efa inventory drift-reconcile history** command.
4. Check the DRC state in detail using the **efa inventory drift-reconcile detail --uuid** command.
5. Confirm the DRC state by comparing with SLX raslogs using the **show logging raslog reverse** command.
6. Check the HTTPS connectivity issues with SLX.

See [HTTPS Connection Issues Between XCO and SLX](#) on page 28.

Drift and Reconcile Resets the Timezone on the SLX

Drift and Reconcile resets the timezone configured on the SLX. The following symptoms are observed:

1. Configure a new timezone on the SLX.
2. Run the reconcile command using EFA or XCO.
3. The timezone on the SLX changes to the one set in EFA or XCO.

```
SLX# show running-config clock
clock timezone Europe/Stockholm

(efa:extreme)extreme@tpvm:~$ efa inventory device timezone list --ip x.x.x.x
+-----+
| IP Address | Timezone | Fabric | AppState |
+-----+
```

```

| x.x.x.x      | Etc/GMT |      | cfg-in-sync |
+-----+
Timezone details
--- Time Elapsed: 117.428303ms ---

(efa:extreme)extreme@tpvm:~$ efa inventory drift-reconcile execute --ip x.x.x.x --reconcile
+-----+
| IP ADDRESS  | RECONCILE | UUID          | STATUS | REASON |
+-----+
| x.x.x.x     | Yes       | <---uuid---> | Success |      |
+-----+

(efa:extreme)extreme@tpvm:~$ efa inventory drift-reconcile history
+-----+
| UUID        | Device IP | Status | Reason | Type | StartTime |
+-----+
| <---uuid---> | x.x.x.x   | success | manual | drift-and-reconcile | 2023-02-09 14:54:11 +0100 CET |
+-----+

SLX# show running-config clock
clock timezone Etc/GMT

```

EFA inventory timezone is set to a different timezone than the one configured on the SLX.

XCO detects this as a drift and reconciles it by configuring the timezone stored in its inventory settings.

Resolution

Configure the new timezone using the following EFA inventory command:

```
efa inventory device timezone set --ip x.x.x.x --timezone Europe/Stockholm
```

Debug Device Settings Changes

When a device setting command fails, an error message will be displayed, providing details on the nature of the failure to facilitate troubleshooting.

```

+-----+
| STATUS | VALUE | ERROR |
+-----+
| Failed | 150   | Configuration of mct bring-up |
|         |       | delay failed for device with |
|         |       | IP: 10.139.44.153 - Device is |
|         |       | not configured on Fabric MCT. |
+-----+

```

Resolution

Complete the following steps:

1. **Compare configurations:** Identify discrepancies between XCO and SLX configurations.
2. **Show device settings:** Use the `efa inventory device setting show -ip <ip>` command to display device settings for a specific IP address.

EFA Firmware-Download Prepare Add Command Fails

The `efa inventory device firmware-download prepare add` command fails:

```
$ efa inventory device firmware-download prepare add --ip 10.147.65.21 --firmware-
host=10.169.71.249 --firmware-directory=/home/cECC_IOS/slxos20.2.3f
Validate Firmware Download Prepare Add [failed]
10.147.65.21: Configured MCT Leaf pair 10.147.65.21, 10.147.65.24 is currently not
connected : missing IP address
--- Time Elapsed: 10.102793812s ---
```

The command failure occurred due to the same ASN numbers being used in two different fabrics.

```
Configured MCT Leaf pair 10.147.65.21, 10.147.65.24 is currently not connected.
```

Resolution

Assign different ASN numbers to each Fabric.

EFA Inventory Config-Replay Execute Command Fails

The `efa inventory config-replay` execution failed after performing backup, deploy, and restore operations.

```
(efa:extreme)extreme-tpvm:~$ efa inventory config-replay detail --uuid
fc3db20b-816a-4ae8-9c3f-97584e535f1d
+-----+-----+
|          NAME          |          VALUE          |
+-----+-----+
| UUID                  | fc3db20b-816a-4ae8-9c3f-97584e535f1d |
+-----+-----+
| Device IP             | 1.1.1.1                 |
+-----+-----+
| Status                | failed                 |
+-----+-----+
| Config Backup UUID    | cfcdb6b9-3849-437a-8c68-e70d216cb07b |
+-----+-----+
| Config Backup SSID    | 1                        |
+-----+-----+
| Failure Reason       | config replay failed  |
+-----+-----+
| Execution Reason      | manual                   |
+-----+-----+
| operation              | manual trigger          |
+-----+-----+
| Start Time            | 2023-05-15 15:56:04 +0200 CEST |
+-----+-----+
| Last Modified Time    | 2023-05-15 15:57:25 +0200 CEST |
+-----+-----+
| Duration              | 1m20.630061514s        |
+-----+-----+
```

The inventory logs indicate that the certificate re-installation could not be found.

Resolution

After restoring the system, ensure to re-install the device certificates.

Execution Error: 503 Service Unavailable

When you run the `efa inventory device firmware-download commit --fabric <fabric name>` command, it responds with `Execution error: 503 Service Unavailable`.

A `503 Service Unavailable Error` is an HTTP response status code indicating that the server is temporarily unable to handle the request. This might occur because the server is overloaded or undergoing maintenance. Unlike a `500 Internal Server Error`, which suggests the server is completely unable to process the request, a `503 Service Unavailable Error` means the server is functioning correctly but cannot handle the request at the moment.

Resolution

This error is usually temporary. Try the command again before assuming the service is down. If the error persists, contact GTAC for further investigation.

HTTPS Connection Issues Between XCO and SLX

Use this topic to learn about troubleshooting HTTPS connection issues between XCO and SLX.

HTTP/HTTPS Process Error and Multiple Defunct Process

The following symptoms are observed when an HTTPS connection from XCO to SLX fails:

- Pings between EFA or XCO and SLX are successful.
- EFA or XCO sends HTTP TCP SYN packets to SLX on port 443.
- SLX does not respond to EFA or XCO. It is found that the HTTP or HTTPS process on SLX has entered a "defunct state."

There are two reasons for HTTPS connection failure from SLX:

- The issue on the SLX device is caused by the HTTP/HTTPS process not responding to the EFA or XCO after receiving a TCP SYN packet on port 443
- Multiple defunct (zombie) processes are present on SLX, indicating that the parent process is unaware of the child process's termination. You cannot kill a zombie process using the SIGKILL signal command. Although these zombie processes are usually inactive and harmless, a large accumulation can cause unexpected behavior.

Resolution

To recover HTTPS connection between EFA or XCO and SLX, complete the following steps:

1. Clear existing HTTPS and OAuth certificates on SLX.
 - List the existing certificates.

```
show crypto ca certificates
```

- Remove OAuth and HTTPS certificates.

```
no crypto import oauth2pkicert
no crypto ca import-pkcs type pkcs12 cert-type https
```

- Verify the removal.

```
show crypto ca certificates
```

2. Restart HTTP service for the Management VRF.

```
http server use-vrf mgmt-vrf shutdownno
http server use-vrf mgmt-vrf shutdown
```

- Ensure that HTTPS is disabled and HTTP is enabled by running the following command:

```
show http server status
```

3. Ensure that there are no certificates and HTTPS is disabled on SLX, but HTTP is enabled by running the following command:

```
show http server status
```

EFA or XCO starts communicating with SLX using HTTP.

4. Configure HTTPS certificates on SLX via EFA or XCO.

```
efa certificate device install --ip10.x.x.46--cert-type token
efa certificate device install --ip10.x.x.46--cert-type http
```

5. If you have removed the syslog CA certificate, reconfigure it.

```
efa inventory device update--ip 10.x.x.46
```

6. Verify Certificates on SLX.

- Check the certificates in the SLX-OS CLI to ensure they are correctly installed.

```
show crypto ca certificates
```

7. Restart HTTP/HTTPS Services.

```
http server use-vrf mgmt-vrf shutdownno
http server use-vrf mgmt-vrf shutdown
```

8. Confirm that HTTPS is enabled.

```
show http server status
```

9. Verify HTTPS Connection.

- Run the following command on EFA or XCO to ensure that the running configuration persist on all fabric devices and SLX is reachable via HTTPS.

```
(efa:extreme)extreme@tpvm21:~$ efa inventory device running-config persist --fabric Pod1
Persist Device(s) Running-Config[success]
```

IP Address	Host Name	Fabric	Status	Reason
10.x.x.45	Pod1-BLeaf2	Pod1	Success	
10.x.x.46	Pod1-Leaf1	Pod1	Success	
10.x.x.47	Pod1-Leaf2	Pod1	Success	

Persist Running-Config Details

--- Time Elapsed: 14.490663218s ---

```
(efa:extreme)extreme@tpvm21:~$ efa fabric show
```

Fabric Name: default, Fabric Description: Default Fabric, Fabric Stage: 3, Fabric Type: clos, Fabric Status: created

IP ADDRESS	POD	HOST NAME	ASN	ROLE	DEVICE STATE	APP
------------	-----	-----------	-----	------	--------------	-----

```
STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
Fabric Name: Pod1, Fabric Description: , Fabric Stage: 3, Fabric Type: clos,
Fabric Status: configure-success
-----+-----+-----+-----+-----+
| 10.x.x.41 | | Pod1-Spine1 | 4200002000 | Spine | | provisioned |
cfg -in-sync | NA | | NA | NA | | 1 |
| 10.x.x.42 | | Pod1-Spine2 | 4200002000 | Spine | | provisioned |
cfg -in-sync | NA | | NA | NA | | 1 |
| 10.x.x.46 | | Pod1-Leaf1 | 4200000000 | Leaf | | provisioned |
cfg -in-sync | NA | | NA | 2 | | 1 |
| 10.x.x.47 | | Pod1-Leaf2 | 4200000000 | Leaf | | provisioned |
cfg -in-sync | NA | | NA | 2 | | 1 |
| 10.x.x.48 | | Pod1-Leaf3 | 4200000001 | Leaf | | provisioned |
cfg -in-sync | NA | | NA | 2 | | 1 |
| 10.x.x.49 | | Pod1-Leaf4 | 4200000001 | Leaf | | provisioned |
cfg-in-sync | NA | | NA | 2 | | 1 |
| 10.x.x.44 | | Pod1-BLeaf1 | 4200003000 | BorderLeaf | | provisioned |
cfg - in-sync| NA | | NA | 2 | | 1 |
| 10.x.x.45 | | Pod1-BLeaf2 | 4200003000 | BorderLeaf | | provisioned |cfg
- in-sync | NA | | NA | 2 | | 1 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```



Note

- For HTTPS to function on SLX, ensure that the appropriate certificates are installed. If no certificate is installed, the HTTPS server will be automatically disabled.
- Configure OAuth and HTTPS certificates.


```
crypto import oauth2pkicert
crypto ca import-pkcs type pkcs12 cert-type http
```
- Enable or disable HTTP or HTTPS.


```
[no] http server use-vrf mgmt-vrf shutdown
```
- If you configure SLX through EFA or XCO, the EFA or XCO will automatically manage the certificate configuration on the SLX during device registration.

Leaf and Spine Nodes Stuck in "cfg-refresh error" After Restore

Following a backup and restore operation on both leaf and spine nodes, and subsequently on the border leaf (BL) nodes, the devices were found in a `cfg-refresh error` state. No configuration changes were made during this process; the devices were simply backed up and restored with the same configuration.

In EFA or XCO, the `config-backup` and `config-replay` processes may sometimes encounter a `cfg-refresh error` usually due to communication issue between EFA or XCO and registered SLX devices. Although this error is uncommon, it can occur. Usually, it resolves on its own once communication is re-established, and the devices will then display a `cfg-refresh error` status. This error may also result from configuration mismatches between the devices and the expected fabric configurations, or from interface flapping on fabric links.

Resolution

1. Wait for automatic drift-reconcile (takes around 1 hour).


```

view | | proto | | proto |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 10.53.17.201 | XCO-USER | GROUPX | view2 |
view2 | view2 | view2 | sha | Bd6GfjluZcVbdtSEeV5tWylqtPg=
| AES128 | Bd6GfjluZcVbdtSEeV5tWylqtPg= | Success | Configured view contains out | | | |
| | | | | | | |
| | | of band configuration. Out of | | | |
| | | | | | | |
| | | band views are not recommended. | | | |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
Snmp user details
--- Time Elapsed: 7.239535431s ---

```

Resolution

- This is a known behavior in the Linux bash prompt. The dollar symbol (\$) is interpreted as a UNIX operator.
- To include a dollar symbol in beginning or middle of the passphrase, use a backslash (\) before the dollar symbol (for example, ab\\$defg123). This suppresses the UNIX operator and treats \$ as a special character in the string. Therefore, enter the resulting passphrase as ('ab\$defg123').



Note

- The issue occurs when the 3rd character in the password is a \$ symbol (for example, ab\$defg123).
- This issue is only present when configured via EFA-CLI, not when configured via SLX-CLI.

- The following is an example of failing Configuration via SLX-CLI & EFA-CLI with differences in encrypted passwords:

```

SLXB(config)# snmp-server user SLX-USER groupname GROUPX auth sha auth-password
ab$defg123 priv AES128 priv-password ab$defg123
(efa:extreme)extreme@tpvm203:~$ efa inventory device snmp user create --ip
10.53.17.201 --name XCO-USER --auth-protocol sha --auth-pass ab$defg123 --priv-
protocol AES128 --priv-pass ab$defg123 --group GROUPX --enable-notify-access --view
view2 --enable-read-access --enable-write-access

```

- The following is the **show run snmp** command output with different encryption keys for the same password:

```

snmp-server user SLX-USER groupname GROUPX auth sha auth-password
O2WqWtstk6iclksgC+qow9wfOvM= priv AES128 priv-password O2WqWtstk6iclksgC+qow9wfOvM=
encrypted
snmp-server user XCO-USER groupname GROUPX auth sha auth-password
Bd6GfjluZcVbdtSEeV5tWylqtPg= priv AES128 priv-password Bd6GfjluZcVbdtSEeV5tWylqtPg=
encrypted

```

Switch Replacement Procedure Fails

While attempting to verify the switch replacement procedure, an error indicated that the system requirements were not met.

```
Checking system configuration on 1.1.1.2 and 1.1.1.3...
System requirements not met on 1.1.1.2.
```

This is a configuration issue due to underscores being disallowed in the hostname.

Resolution

Reconfigure the hostname to exclude underscores.

Verify XCO is Registered to Receive Traps from SLX

Verify the following configuration to confirm that XCO is registered to receive traps from SLX:

- Verify the following configuration on SLX device:

```
SLX# show run snmp
snmp-server sys-descr "Extreme SLX9250 Switch"
snmp-server user efav3User groupname efav3Group
auth sha auth-password
"LlgMJxyldZfAfPydPDdFTNNtkpw=\n" priv AES128 privpassword
"53IJSJKiK/A6UFLVd9+ZkkclIJs=\n" encrypted
snmp-server v3host 10.20.241.85 efav3User
source-interface management chassis-ip
!
snmp-server view efav3View 1.3.6.1 included
snmp-server view v3view 1.3.6.1 included
snmp-server group efav3Group v3 read efav3Group
write efav3Group notify efav3View
snmp-server group v3group v3 read v3group write
v3group notify v3view
```

- Another verification method is to set up a trap forwarding recipient and then confirm if the receiver receives the traps.



Note

SNMP service logs are available at `/var/log/efa/snmp/snmpserver.log`.

Verify the Types of Device Updates Done

Use this topic to learn about structured messages containing necessary details for both the minimal and full updates, including timestamps, application names, device identifiers, use cases, log levels, messages, and request IDs.

The example messages clearly indicate the type of update being performed, the reason for the update, and other relevant details.

Minimal Update Message

To create a minimal update message with the reason of `RaslogEventReason` for the device 10.139.44.175, structure the message as follows.

```
{
  "@time": "2022-04-12T16:05:54.242995 EDT",
  "App": "inventory",
  "Device": "10.139.44.175",
  "UseCase": "update Devices",
  "level": "info",
  "msg": "ComputeDeviceUpdates: Starting minimal update with reason(RaslogEventReason)",
  "rqId": "fce6e236-1fe8-4d00-bf7a-4fc8d057058c"
}
```

Full Update Message

To create a full update message indicating that since “all:true” is set, everything on the device will be queried, structure the message as follows.

```
{
  "@time": "2022-04-12T18:37:04.219736 EDT",
  "App": "inventory",
  "Device": "10.139.44.175",
  "UseCase": "update
Devices", "level": "info",
  "msg": "ComputeDeviceUpdates: Starting full update to do TriggerDeepDeviceUpdate with
reason(ManualUpdate) mapEvts(map[all:true])",
  "rqId": "e4e4f04b-9913-4d11-85e2-acc9b96face6"
}
```

If the query was only for the interface information “Interface:true”, then only interface info is queried from the device.



Note

To get the current state of devices or to perform on-demand full device discovery, run the **efa inventory device update** command.

Verify the Device Supports Minimal Updates

Ensure that the SLX version is 20.4.3a or greater.

If the following log shows that the `configDriftSupport` is set to true, then the device supports tracking out-of-band configuration changes, which XCO can use to facilitate minimal updates.

```
{
  "@time": "2022-04-16T12:07:04.830832 EDT",
  "App": "inventory",
  "Device": "10.139.44.175",
  "UseCase": "update Devices",
  "level": "info",
  "msg": "ComputeDeviceUpdates: reason(RaslogEventReason) configDriftSupport(true)
oldCounter(17) oldTimestamp(2022-04-14 16:11:56 -0400 EDT) to newCounter(18)
newTimestamp(2022-04-16 12:04:27 -0400 EDT)",
  "rqId": "ac13a3c9-c0ae-4317-a16c-b9f055007712"
}
```

Verify Device Update Notification on Out of Band Config Change on SLX

In the following log, if the `oldCounter` and `newCounter` values differ, XCO has determined that an out of band configuration change was made either using either the `efa inventory device execute-cli` command or other non-XCO mechanisms via the SLX CLI, REST or Netconf.

```
{
  "@time": "2022-04-16T12:07:04.830832 EDT",
  "App": "inventory",
  "Device": "10.139.44.175",
  "UseCase": "update Devices",
  "level": "info",
  "msg": "ComputeDeviceUpdates: reason(RaslogEventReason) configDriftSupport(true)
oldCounter(17) oldTimestamp(2022-04-14 16:11:56 -0400 EDT) to newCounter(18)
newTimestamp(2022-04-16 12:04:27 -0400 EDT)",
  "rqId": "ac13a3c9-c0ae-4317-a16c-b9f055007712"
}
```

Verify Out of Band Config Change (Non-XCO) Triggered XCO for a Full Update

In the following log, "all:true" is added to the map of events to be updated because XCO determined there was an out-of-band (OOB) configuration change by comparing the `oldCounter` and `newCounter` values.

Initially, the update was set to focus only on interfaces with "Interface." However, upon detecting the OOB changes, an event was added to update everything with "all:true".

```
{
  "@time": "2022-04-16T12:07:04.832388 EDT",
  "App": "inventory",
  "Device": "10.139.44.175",
  "UseCase": "update Devices",
  "level": "info",
  "msg": "ComputeDeviceUpdates: config drift counter has changed so a full update needs
to be done mapEvts(map[Interface:true all:true])",
  "rqId": "ac13a3c9-c0ae-4317-a16c-b9f055007712"
}
```

Verify Reachability of XCO to a New Management IP Address

To confirm that the XCO server is reachable from a new management IP address, make sure the VLAN configuration of the sub-interface and the IP address subnet are compatible with the host attempting to establish a connection (ping) to the XCO server.

Verify Switch Registration Process Completion on XCO Using REST APIs

Using the XCO REST APIs, you can check the completion status of the switch registration process, including SNMP view creation.


To determine if the registration process is complete, use the `GetSwitches` option under inventory with the endpoint `get inventory/switches?device_ips=x.x.x.x`.

Before you configure anything on the SLX, ensure that the `admin_state` is "up", the `device_health` is "healthy," and the `discovery_status` is "completed".

```
2023-06-28T04:28:28.580Z
SENT: get https://x.x.x.y/v1/inventory/switches?device_ips=x.x.x.x

2023-06-28T04:28:28.682Z
RECEIVED: 200:{"items":[{"fabric":
{"ip_address":"x.x.x.x","mac_address":"xx:xx:xx:xx:xx:xx","location":"default","name":"S
LX","id":179,"device_type_id":3,"device_type":"SLX","model":"3200","chassis_name":"8720-32
C","firmware":"20.5.1a","type":"FABRIC","discovery_status":"completed","global_l2_mtu":921
6,"global_ip_mtu":1500,"last_discovery_time":"0001-01-01T00:00:00Z","discovery_interval":"
0","admin_state":"up","device_health":"healthy","is_maintenance_mode_enabled":false,"is_ma
intenance_mode_on_reboot_enabled":false,"is_syslog_configured":true,"switch_global_config_
data":
{"mac_aging_timeout":1800,"conversational_mac_aging_timeout":300,"conversational_move_limi
t":20}}]}
```

For more information, see [getSwitches](#).



Troubleshooting Fabric Skills

- [Access Token Expires Before the Specified Expiration Time](#) on page 38
- [Active XCO or EFA Node is Down](#) on page 39
- [Cannot Ping IP Address of Sub-Interface](#) on page 40
- [Check Status of Restore API](#) on page 40
- [Daemonset "goraslog-service" is Stuck in a "Not Ready" State](#) on page 40
- [Devices in CFG Refreshed or CFG Refreshed Error State](#) on page 41
- [Device Addition to Fabric Fails](#) on page 49
- [EFA Status Shows Nothing and gofaultmanager and gopolicy POD Services Stuck in init State](#) on page 50
- [EFA Commands Fail to Execute](#) on page 51
- [EFA Commands Fails with "Error 408"](#) on page 52
- [EFA Command Fails with "Dial TCP xx.xx.xx.xx:80: Connect: Connection Refused" Message](#) on page 52
- [EFA Error When Registering a RELP Handler on Endpoint Using FQDN](#) on page 52
- [EFA Fabric Show Command Fails After Active Node Restart Causing CFG Refreshed Error on Devices](#) on page 53
- [EFA or XCO is Down on Standby TPVM](#) on page 54
- [EFA Fabric health Not in Sync with EFA Fabric Topology Show](#) on page 56
- [EFA or XCO VM is down and Came Up After Stopping the Services](#) on page 57
- [EFA or XCO Execute-CLI Reports SLX Switches as Unreachable](#) on page 57
- [EFA or XCO Execute-CLI Fails on SLX](#) on page 58
- [EFA or XCO Fails to Resolve Non-Cluster FQDN](#) on page 58
- [EFA or XCO Fails to Log in to Switch](#) on page 59
- [Fabric Skill Troubleshooting](#) on page 60
- [Fabric Skill is Down After a Failed Token Certificate Renewal](#) on page 62
- [Fabric Configuration Fails with EFA or XCO](#) on page 65
- [How to Add a Second Link to an Existing Fabric via EFA or XCO](#) on page 66
- [Log in to EFA or XCO Fails](#) on page 66
- [LDAP Authentication Configuration Using FQDN Fails](#) on page 67
- [Management Subinterface IP Address Not Listed](#) on page 67
- [Recover a Leaf Node after Deletion](#) on page 68
- [SLX, EFA, or XCO Config Missing After SLX Reload](#) on page 69
- [TPVM Unable to Authenticate to Red Hat Directory Service \(LDAP\)](#) on page 70

- [Unable to Run EFA Commands on page 70](#)
- [Unable to Connect to the Server on page 71](#)
- [Update Authentication Preference on page 72](#)
- [Update of Maximum Password Age of SLX Password Fails on page 73](#)
- [Analyze 10 Second Traffic Loss During DN Leaf Power Cycle on page 73](#)
- [XCO or EFA Authentication Stops Working when Users Configure a Different Authentication on page 74](#)
- [EFA CLI Login Fails with Error "Failed to Fire Hook" on page 76](#)

Learn about troubleshooting fabric skills issues.

Access Token Expires Before the Specified Expiration Time

Access token expiration set to 1 hour, but REST commands fail after a few minutes.

```
INFO:root:Token expiry time: {'type': 'ACCESS', 'hours': 1}
INFO:root:EFA token expiry updated (2021-02-15 15:25:41.578542)
INFO:xxxxxxxx - - [15/Feb/2021 14:27:21] "GET /consumer HTTP/1.1" 200 -
INFO:xxxxxxxx - - [15/Feb/2021 14:28:11] "GET /consumer HTTP/1.1" 200 -
INFO:xxxxxxxx - - [15/Feb/2021 14:29:11] "GET /consumer HTTP/1.1" 200 -
INFO:xxxxxxxx - - [15/Feb/2021 14:30:11] "GET /consumer HTTP/1.1" 200 -
INFO:xxxxxxxx - - [15/Feb/2021 14:31:08] "GET /consumer HTTP/1.1" 200 -
INFO:xxxxxxxx - - [15/Feb/2021 14:31:09] "GET /consumer HTTP/1.1" 200 -
INFO:xxxxxxxx - - [15/Feb/2021 14:31:09] "POST /consumer HTTP/1.1" 401 -
INFO:root:Unable to create EFA Tenant (token is expired by 2m58s)
```

Updating the expiration to 24 hours, 12 hours, and 1 hour yields the same behavior.

```
{ "App": "efa", "level": "info", "msg": "Completed", "reason": "Token Expiry Time Has Been Successfully Updated", "request": { "cmd": "auth:update-token-expiry", "params": { "Type": "ACCESS" }, "rqId": "2523377a-7f0c-424c-a034-fa67959881f4", "time": "2021-02-15T14:28:43Z" } }
{"level": "info", "msg": "New request to validate token", "time": "2021-02-15T14:28:51Z"}
{"level": "info", "msg": "Token is validated successfully", "time": "2021-02-15T14:28:51Z"}
{"level": "info", "msg": "Added user header 'admin' for the request", "time": "2021-02-15T14:28:51Z"}
{"level": "info", "msg": "Added user role 'SystemAdmin' for the request", "time": "2021-02-15T14:28:51Z"}
{"level": "info", "msg": "Roles in token context: [SystemAdmin]", "time": "2021-02-15T14:28:51Z"}
{"level": "info", "msg": "Request is not ML2 request", "time": "2021-02-15T14:28:51Z"}
{"App": "efa", "level": "info", "msg": "Request ID Created", "rqId": "537f4a93-c41d-4bb8-9196-d28f52a059af", "time": "2021-02-15T14:28:51Z"}
{"App": "efa", "level": "info", "msg": "Recieved", "request": { "cmd": "auth:get-token-expiry", "params": { "Type": "ACCESS" }, "rqId": "537f4a93-c41d-4bb8-9196-d28f52a059af", "time": "2021-02-15T14:28:51Z"}
{"App": "efa", "level": "info", "msg": "Completed", "reason": "Fetched token expiry settings successfully.", "request": { "cmd": "auth:get-token-expiry", "params": { "Type": "ACCESS" }, "rqId": "537f4a93-c41d-4bb8-9196-d28f52a059af", "time": "2021-02-15T14:28:51Z"}
{"level": "info", "msg": "New request to validate token", "time": "2021-02-15T14:28:52Z"}
{"level": "info", "msg": "Token is validated successfully", "time": "2021-02-15T14:28:52Z"}
{"level": "info", "msg": "Added user header 'admin' for the request", "time": "2021-02-15T14:28:52Z"}
{"level": "info", "msg": "Added user role 'SystemAdmin' for the request", "time": "2021-02-15T14:28:52Z"}
{"level": "info", "msg": "Roles in token context: [SystemAdmin]", "time": "2021-02-15T14:28:52Z"}
```

```
{ "level": "info", "msg": "Request is not ML2 request", "time": "2021-02-15T14:28:52Z" }
{ "level": "info", "msg": "New request to validate token", "time": "2021-02-15T14:29:42Z" }
{ "level": "error", "msg": "invalid authorization header: token
is expired by 1s", "time": "2021-02-15T14:29:42Z" }
```

The issue occurs when you update the token expiry settings after the token was created.

Resolution

Token expiry settings apply at the time of token creation.

Updating the token expiration to one hour means that all tokens generated after this update will expire in one hour. Existing tokens are not affected by this change.

Active XCO or EFA Node is Down

When you run the **show efa status** command from the SLX command prompt, the command output show that active XCO or EFA node is down.

```
SLX# show efa status
=====
EFA version details
=====
Version : 3.2.1
Build: GA
Time Stamp: 23-05-16:02:55:28
Mode: Secure
Deployment Type: multi-node
Deployment Platform: TPVM
Deployment Suite: Fabric Automation
Deployment IP Mode: ipv4
Virtual IP: x.x.x.x
Node IPs: y.y.y.y,z.z.z.z
--- Time Elapsed: 42.26278ms ---

=====
EFA Status
=====
+-----+-----+-----+-----+
| Node Name | Role   | Status | IP      |
+-----+-----+-----+-----+
| el2c1tpvm2 | active | down   | y.y.y.y |
+-----+-----+-----+-----+
| el1c1tpvm1 | standby | up     | z.z.z.z |
+-----+-----+-----+-----+
```

Resolution

Verify if the RabbitMQ logs are rotating in XCO.

- If they are, upgrade to XCO 3.3.1 or above where the log rotation script is set up automatically.

For more information, see [RabbitMQ Logs Do Not Rotate in XCO 3.2.x and Above](#).

Alternatively, complete the following workaround:

1. Truncate the log file or clean all the logs using the **efactl clean** command.
2. Restart the EFA or XCO service using the **efactl restart** command.
3. If any pods are in the evicted state, reimport the images using the **k3s ctr image import /opt/efa/docker_images/docker_k3s_images.tar** command:
4. Restart the k3s service using the **systemctl restart k3s** command.

Cannot Ping IP Address of Sub-Interface

Ping to the IP address of the sub-interface created from the host which is on the same subnet fails.

Follow this procedure to ping IP address of a sub-interface.

1. On an active node, run the **ip addr show type vlan** command.
 - a. Check the vlan/subinterface information for IP address.
 - b. Check if the interfaces are up and running.
2. Check the keepalived configuration at the following location:
/etc/keepalived/keepalived.conf

Look for IP address of sub interfaces in `vrrp_instance` under `virtual_ipaddress`.

3. Check if the traefik or ingress policy is configured correctly using the **k3s kubectl get service traefik -n kube-system** command.

Check the external IP for the sub interface IP addresses and the TCP ports 80 and 443.

Check Status of Restore API

The restore API is updated to an async API. It returns a 202 response with an ID of the restore operation in progress. You can use the ID to check the status of the operation by using the endpoint `/v1/system/restore/{id}`.

System restore detailed logs are available at the `<Log directory>/system/system-client.log` location.

Daemonset "goraslog-service" is Stuck in a "Not Ready" State

During fabric installation, daemonset "goraslog-service" is stuck in "not ready" state.

AME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE	NODE
SELECTOR		IMAGES		SELECTOR		
daemonset.apps/goopenstack-service	0	0	0	0	0	non-
existing=true	209d	openstack	goopenstack:3.0.1	app=goopenstack-service		
daemonset.apps/govcenter-service	0	0	0	0	0	non-
existing=true	209d	vcenter	govcenter:3.0.1	app=govcenter-service		
daemonset.apps/gohyperv-service	0	0	0	0	0	non-
existing=true	209d	hyperv	gohyperv:3.0.1	app=gohyperv-service		
daemonset.apps/goraslog-service	0	0	0	0	0	non-
existing=true	209d	goraslog-service	goraslog:3.0.1	app=goraslog-service		

daemonset.apps/gosnmp-service	1	1	1	1	1
<none>	209d	gosnmp-service	gosnmp:3.0.1	app=gosnmp-service	
daemonset.apps/gosystem-service	1	1	1	1	1
<none>	209d	gosystem	gosystem:3.0.1	app=gosystem-service	
daemonset.apps/efa-api-docs	1	1	1	1	1
<none>	209d	efa-api-docs	efa-api-docs:3.0.1	app=efa-api-docs	
daemonset.apps/rabbitmq	1	1	1	1	1
<none>	209d	rabbitmq-node	rabbitmq:3.0.1	app=rabbitmq	
daemonset.apps/gofabric-service	1	1	1	1	1
<none>	209d	gofabric-service	gofabric:3.0.1	app=gofabric-service	
daemonset.apps/gorbac-service	1	1	1	1	1
<none>	209d	go-rbac	gorbac:3.0.1	app=gorbac-service	
daemonset.apps/goinventory-service	1	1	1	1	1
<none>	209d	goinventory-service	goinventory:3.0.1	app=goinventory-service	
daemonset.apps/gonotification-service	1	1	1	1	1
<none>	209d	gonotification-service	gonotification:3.0.1	app=gonotification-service	
daemonset.apps/goauth-service	1	1	1	1	1
<none>	209d	go-auth	goauth:3.0.1	app=goauth-service	
daemonset.apps/gopolicy-service	1	1	1	1	1
<none>	209d	gopolicy-service	gopolicy:3.0.1	app=gopolicy-service	
daemonset.apps/gotenant-service	1	1	1	1	1
<none>	209d	gotenant-service	gotenant:3.0.1	app=gotenant-service	

It was observed during a planned reboot of BL01 the "exit MM" did not occur. Subsequently, it was noted that daemonset.apps/goraslog-service was not ready, and the last RASLOG entries were from August 2nd (EFA - raslog/raslog-server.log).

Resolution

Manually restart the service using **sudo efactl start-service goraslog** command.

```
(efa:extreme)extreme@tpvm2:/apps/efa_logs/rabbitmq$ sudo efactl start-service goraslog
Are you sure you want to start "goraslog"? [Y/n]
Y
"goraslog" has been started
```

Devices in CFG Refreshed or CFG Refreshed Error State

Use this topic to learn about the troubleshooting scenario when a device is in `cfg refreshed` state.

EFA or XCO Device Status is "cfg refreshed" or "cfg refresh error"

When running the **efa fabric show** command, some devices are listed as `cfg refreshed` or `cfg refresh error`. This indicates a discrepancy between the XCO or EFA configuration and the actual device configuration.

```
(efa:extreme)extreme@tpvm21:~$ efa fabric show

Fabric Name: extr_pdl, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric Status:
configure-success, Fabric Health: Black
+-----+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE | DEVICE |
STATE | APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
```

```

| 10.53.2.41 | | PD1_S1 | 4200002000 | Spine | provisioned
| cfg in-sync | NA | NA |
| 10.53.2.42 | | PD1_S2 | 4200002000 | Spine | provisioned
| cfg in-sync | NA | NA |
| 10.53.2.46 | | PD1_L01 | 4200000000 | Leaf | provisioned
| cfg in-sync | NA | NA |
| 10.53.2.47 | | PD1_L02 | 4200000000 | Leaf | provisioned
| cfg refreshed | LA,LD,IU,BGPU | BGP-D,INTIP-C,INTIP-U,INTIP-D | NA | 1 |
| 10.53.2.44 | | PD1_L03 | 4200003000 | BorderLeaf | provisioned
| cfg refresh error | LA,ASN,IU,POU | MCT-C,MCT-PA,INTIP-C,EVPN-C,O-C | 2 | 1 |
| 10.53.2.45 | | PD1_BL2 | 4200003000 | BorderLeaf | provisioned
| cfg refresh error | LA,LD,ASN,IU,POU | MCT-C,MCT-PA,BGP-U,INTIP-U | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The issue occurred due to the conflicts between EFA or XCO and device configuration, and due to the various situations, such as networks updates, BGP changes, or system properties updates.

Resolution

Use either of the following methods to overcome the `cfg refreshed` or `cfg refresh error`:

- Revert SLX configuration to match XCO config.
Update inventory using the `efa inventory device update --ip [ip_address]` command. This will confirm whether the SLX config is back in sync.
- Perform Drift Reconcile (DRC).

Run the `efa inventory drift-reconcile execute --ip [ip_address] --reconcile` command to overwrite conflicting config.

Optionally, run without the `--reconcile` flag to display a summary of conflicting configs.

EFA Still Reporting "cfg refreshed" Post Drift Reconcile

Despite a successful drift reconciliation, EFA or XCO continues to report `cfg refreshed` for several interfaces on the device. The issue persists even after running the `efa tenant debug device drift` command with the `--reconcile` option.

The command output shows that the bridge-domain and lif drift configurations were successfully reconciled, but the EFA or XCO still reports a drifted config for the affected interfaces.

```

efa:admin)admin@efadeployment[/home/admin]$efa tenant debug device drift --device-ip 10.158.8.65 --
reconcile
=====
=====
Device           : 10.158.8.65
===== BD Drift =====
BD               : 12
App-state        : cfg-in-sync

Drifted Interface
+-----+-----+-----+-----+-----+-----+-----+-----+

```

```

| Interface-Type | Interface-Name | App-State |
+-----+-----+-----+
| ethernet      | 0/5.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/4.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/3.3          | cfg-refreshed |
+-----+-----+-----+
===== Lif Drift =====

Drifted Lif
+-----+-----+-----+
| Interface-Type | Interface-Name | App-State |
+-----+-----+-----+
| ethernet      | 0/5.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/4.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/3.3          | cfg-refreshed |
+-----+-----+-----+

===== Reconciliation Status =====
+-----+-----+-----+
| CONFIG TYPE | STATUS | ERROR |
+-----+-----+-----+
| Bridge-domain | Success | |
| Lif           | Success | |
+-----+-----+-----+

Operation succeeded.

Then checked again the EFA still report drifted config:

efa:admin)admin@efadeployment[/home/admin]$date
Thu Jun 22 09:45:00 CEST 2023
efa:admin)admin@efadeployment[/home/admin]$efa tenant debug device drift --device-ip 10.158.8.65
=====
Device           : 10.158.8.65
===== BD Drift =====
BD               : 12
App-state        : cfg-in-sync

Drifted Interface
+-----+-----+-----+
| Interface-Type | Interface-Name | App-State |
+-----+-----+-----+
| ethernet      | 0/5.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/4.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/3.3          | cfg-refreshed |
+-----+-----+-----+
===== Lif Drift =====

Drifted Lif
+-----+-----+-----+
| Interface-Type | Interface-Name | App-State |
+-----+-----+-----+
| ethernet      | 0/5.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/4.3          | cfg-refreshed |
+-----+-----+-----+
| ethernet      | 0/3.3          | cfg-refreshed |
+-----+-----+-----+

```

```
+-----+  
Operation succeeded.
```

Resolution

Upgrade to XCO 3.3.0 or above.

"Error: dial tcp" When Running "efa fabric show" Command

After restarting the active node, the **efa fabric show** command returns a `dial tcp` error, and some devices show a `cfg refresh error` status.

The following symptoms are observed:

- When MariaDB is restarted on the active XCO node, the **efa fabric show** command returns an error message indicating a refused connection and invalid transaction.

```
"Error : dial tcp <xco-ip>:3306: connect: connection refused; invalid transaction;  
invalid transaction"
```

- After some time, the database connection is reestablished, but one of the border leaf devices shows a `cfg-refresh-error` status.

The **efa fabric show** command returns a list of devices with their statuses, including the border leaf device showing the error.

```
(efa:extreme)extreme@slx-tpvm:$ efa fabric show  
  
| 192.168.246.21 |      | SPINE01      | 64512 | Spine      | provisioned | cfg refreshed |  
LA,LD,IU      | BGP-D,INTIP-C,INTIP-D      | NA | 1 |  
| 192.168.246.22 |      | SPINE02      | 64512 | Spine      | provisioned | cfg refreshed |  
LA,LD,IU,BGPU | BGP-D,INTIP-C,INTIP-U,INTIP-D | NA | 1 |  
| 192.168.246.31 |      | LEAF001      | 65000 | Leaf       | provisioned | cfg in-sync   |  
NA            | NA            |      | 2 | 1 |  
| 192.168.246.32 |      | LEAF002      | 65000 | Leaf       | provisioned | cfg in-sync   |  
NA            | NA            |      | 2 | 1 |  
| 192.168.246.33 |      | LEAF003      | 65001 | Leaf       | provisioned | cfg in-sync   |  
NA            | NA            |      | 2 | 1 |  
| 192.168.246.34 |      | LEAF004      | 65001 | Leaf       | provisioned | cfg refreshed |  
LA,LD,IU,POU  | MCT-U,MCT-PA,BGP-D,INTIP-C,INTIP-U,INTIP-D | 2 | 1 |  
| 192.168.246.35 |      | BLEAF01      | 66000 | BorderLeaf | provisioned | cfg refresh error |  
LA,LD,IU,POU  | BGP-D,INTIP-C,INTIP-D      | 2 | 1 |  
| 192.168.246.36 |      | BLEAF02      | 66000 | BorderLeaf | provisioned | cfg refreshed |  
LA,LD,IU      | BGP-D,INTIP-C,INTIP-D      | 2 | 1 |  
+-----+  
+-----+
```

The issue occurred because the Inventory database was not fully updated with device status after it came back up.

Resolution

To update the app-state for the device, the complete the following steps:

1. Shut down MCT ports on SLX.
2. Update the device inventory on XCO using the **efa inventory device update -ip <device-ip>** command.

3. Re-enable MCT ports on SLX.
4. Update the device inventory on XCO again using the same command.
5. If the issue persists, perform a Drift Reconcile on the affected device(s) using the **efa inventory drift-reconcile execute --ip <device-ip> --reconcile** command.

Device in "cfg refresh error" After MCT CCP Restoration

The SLX app-state gets stuck in `cfg refresh error` after restoring MCT CCP, toggling MCT cluster ports, experiencing a fiber cut, or reloading a MCT peer.

It is observed that the SLX app-state remains in `cfg refresh error` for an extended period.

The **efa fabric show --name B168-FABRIC** command returns a list of devices with their app-states, including some in `cfg in-sync` and others in `cfg refreshed` or `cfg refresh error`.

This is an expected behavior and a transient state. The app-state will update to `cfg-sync` when the next DRC (Drift Reconciliation) occurs.

Resolution

If the `cfg refreshed` state persists for over 30 minutes or causes functional issues, complete the following steps:

1. Run the **efa inventory device update --ip [SLX_IP]** command.
2. Run the **efa inventory drift-reconcile execute --ip <device-ip> --reconcile** command.

Multiple or Duplicate IPs on Fabric Interfaces

When running the **efa fabric configure** command, duplicate IP addresses appear on fabric interfaces, causing a brief outage between Spine and Border Leaf (BL) devices.

The following symptoms are observed:

- Multiple IP addresses assigned to fabric interfaces connecting Spine to BL
- Original (Brownfield) IP addresses ignored, and new IP addresses allocated.

The issue occurred because there is a out-of-order sequence of events during the configuration process, resulting in incorrect IP address allocation.

Running the `GetInterfaceConfigs` shows new IP addresses assigned for point-to-point links.

Resolution

Upgrade to XCO 3.4.0 or above. Alternatively, complete the following workaround:

1. Shutdown both interfaces via SLX-OS.
2. Remove new IP pairs from both interfaces using the `no ip address [ip_address/cidr]` command.
3. Trigger inventory device update from XCO with `efa inventory device update --ip [switch_ip]` command.
4. Restore both interfaces using the `no shutdown` command.
5. Trigger inventory device update from XCO again.
6. Run the `efa fabric configure --name [fabric_name]` command.

App-state Showing 'cfg-refreshed' for OOB L2VPN

XCO fabric devices display `cfg refreshed` status due to missing 'encapsulation vxlan' and 'enable-peer-as-check' configurations for Out-of-Band (OOB) L2VPN.

The issue occurs when these configurations are absent from the SLX devices, despite being present in the running-config for the XCO-configured "Spine-group" peers. This is because XCO sorts peer group entries by the lowest peer IP address and doesn't verify the correct peer group configured by XCO.

Resolution

Upgrade to XCO 3.6.0 or above.

Device Showing "cfg refreshed" After Annual Interface Description Update

After manually updating the interface description of an ICL (Inter-Chassis Link) on a Leaf2 device from `clusterPeerIntfMember` to `Link1TOPeer` via SLX-OS CLI, the XCO fabric device shows a `cfg refreshed` status.

Log messages indicate a drift in the MCT (Multi-Chassis Trunk) configuration, specifically, in the interface description. The device is flagged for a drift reconcile.

```
{"@time":"2024-03-25T12:39:17.786193 EDT","level":"info","msg":"DriftCompare: identifyFabricMCTDrift Refreshed MCT cluster Diff F|A: Description: clusterPeerIntfMember | Link1TOPeer "}
{"@time":"2024-03-25T12:39:17.786326 EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep Discovery","level":"info","msg":"Check LACP timeout and Description drift for MCT linked interface : {0 0 0 19 ethernet 0/10 40Gbps Ethernet 0/10 00:04:96:d7:04:d4 edge short clusterPeerIntfMember 64 active standard 0 0 0 false }","rqId":"f3f8bad3-6dbe-4d63-96a2-c2c9e9a382c3"}
{"@time":"2024-03-25T12:39:17.788932 EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep Discovery","level":"info","msg":"Is Drift present in MCT: true","rqId":"f3f8bad3-6dbe-4d63-96a2-c2c9e9a382c3"}
{"@time":"2024-03-25T12:39:17.789022 EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep Discovery","level":"info","msg":"Cluster drift is : {AppState:0x2b87cc0 Name:Cluster ChildConfig: [MCTInterfaceDescription:0/9]} \n ","rqId":"f3f8bad3-6dbe-4d63-96a2-c2c9e9a382c3"}
{"@time":"2024-03-25T12:39:17.789108 EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep Discovery","level":"info","msg":"Is Drift present : true","rqId":"f3f8bad3-6dbe-4d63-96a2-c2c9e9a382c3"}
```

XCO manages interface descriptions, and the manual change caused a discrepancy between the expected and actual descriptions, resulting in the `cfg refreshed` status.

Resolution

Perform a Drift Reconcile on the device using the `efa inventory drift-reconcile execute --ip [ip_address] --reconcile` command to correct the configuration.

Device in "cfg refresh error" After BMC Upgrade

After rebooting the spine, some leaves in the XCO IP fabric are experiencing a `cfg refresh error`. The affected leaves are showing the following status:

```
Fabric Name: Fabric_XXX, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric Status:
configure-success, Fabric Health: Black
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE | DEVICE STATE | APP STATE |
CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.10.100.111 | | SLX-Leaf-XXXXXXXX-01 | 65019 | Leaf | provisioned | cfg refresh error |
LD,IU | BGP-D,INTIP-D | 2 | 1 |
| 10.10.100.112 | | SLX-Leaf-XXXXXXXX-02 | 65009 | Leaf | provisioned | cfg refresh error |
LD,IU | BGP-D,INTIP-C,INTIP-D | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

The issue is caused by a device connected to the leaf having an LLDP chassis ID that exceeds the 50-character limit supported by XCO. This results in a bulk LLDP update rejection and LLDP mismatch.

Resolution

Change the connected device's Chassis ID to less than 50 characters.

EFA Fabric Show Command Output Showing Nodes with "cfg refreshed" App State

The output of the `efa fabric show` command shows various nodes with different app states, including "cfg refreshed", "cfg in-sync", and "cfg refresh error":

```
(efa:extreme)extreme@val90-tpvml:~$ efa fabric show

xxxtruncatedxxxxx
Fabric Name: VXQUSDL-fabric-01, Fabric Description: , Fabric Stage: 3, Fabric Type: clos, Fabric
Status: settings-updated, Fabric Health: Black

Updated Fabric Settings: BGP-MD5

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | POD | HOST NAME | ASN | ROLE |
| DEVICE STATE | APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 192.168.246.101 | | VXQUSDS01 | 64512 | Spine |
provisioned | cfg in-sync | NA | NA | NA | 1 |
| 192.168.246.102 | | VXQUSDS02 | 64512 | Spine |
provisioned | cfg refresh error | NA | INTIP-C,INTIP-D | NA | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

```

| 192.168.246.21 | | VXQUSDL01 | 65000 | Leaf |
provisioned | cfg refreshed | LA,LD,IU,POU | MCT-PD,INTIP-C | 2 | 1 |
| 192.168.246.22 | | VXQUSDL02 | 65000 | Leaf |
provisioned | cfg refreshed | LA,LD,IU,POU | MCT-PD,INTIP-C | 2 | 1 |
| 192.168.246.23 | | VXQUSDL03 | 65001 | Leaf |
provisioned | cfg in-sync | NA | NA | 2 | 1 |
| 192.168.246.24 | | VXQUSDL04 | 65001 | Leaf |
provisioned | cfg in-sync | NA | NA | 2 | 1 |
| 192.168.246.25 | | VXQUSDB01 | 66000 | BorderLeaf |
provisioned | cfg in-sync | NA | NA | 2 | 1 |
| 192.168.246.26 | | VXQUSDB02 | 66000 | BorderLeaf |
provisioned | cfg refreshed | LA,LD,IU | BGP-D,INTIP-C | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
FABRIC SETTING:
    
```

The issue arises due to changes in the network configuration, such as moving cables and changing MCT ports.

Resolution

Run the following commands:

- Run the `efa inventory device update --ip [ip_address]` command 2-3 times.
- Run the `efa fabric configure --name <fabric-name>` command.



Note

Error messages indicate issues with IP address overlaps and netconf RPC errors. The fabric drift-reconcile operation failed due to these errors.

XCO Fabric Device Displays "cfg refreshed" After Manual Update of Interface Description

The following symptoms are observed:

- Manually updated the interface description of the ICL from 'clusterPeerIntfMember' to 'Link1TOPeer' on Leaf2 via SLX-OS CLI.
- XCO shows the device in the `cfg refreshed` app-state.
- The log displays the following information:

```

{"@time":"2024-03-25T12:39:17.786193 EDT","level":"info","msg":"DriftCompare:
identifyFabricMCTDrift Refreshed MCT cluster Diff F|A: Description: clusterPeerIntfMember |
Link1TOPeer "}
{"@time":"2024-03-25T12:39:17.786326
EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep
Discovery","level":"info","msg":"Check LACP timeout and Description drift for MCT linked
interface : {0 0 0 19 ethernet 0/10 40Gbps Ethernet 0/10 00:04:96:d7:04:d4 edge
short clusterPeerIntfMember 64 active standard 0 0 0 false }","rqId":"f3f8bad3-6dbe-4d63-96a2-
c2c9e9a382c3"}
{"@time":"2024-03-25T12:39:17.788932
EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep
Discovery","level":"info","msg":"Is Drift present in MCT: true","rqId":"f3f8bad3-6dbe-4d63-96a2-
c2c9e9a382c3"}
{"@time":"2024-03-25T12:39:17.789022
EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep
Discovery","level":"info","msg":"Cluster drift is : {AppState:0x2b87cc0 Name:Cluster ChildConfig:
[MCTInterfaceDescription:0/9]} \n ","rqId":"f3f8bad3-6dbe-4d63-96a2-c2c9e9a382c3"}
{"@time":"2024-03-25T12:39:17.789108
EDT","App":"dcfabric","Device":"10.64.208.28","Fabric":"fabric1","UseCase":"Device Deep
    
```

```
Discovery", "level": "info", "msg": "Is Drift present : true", "rqId": "f3f8bad3-6dbe-4d63-96a2-c2c9e9a382c3"}

```

- DriftCompare identified a drift in the MCT cluster description.
- Device Deep Discovery detected a drift in LACP timeout and description for the MCT linked interface.
- Drift is present in MCT and cluster drift is detected.

Interface descriptions are managed by XCO, and the manual update caused a drift.

Resolution

Perform a Drift Reconcile on the device using the **efa inventory drift-reconcile execute --ip [ip_address] --reconcile** command to correct the config.

Device Addition to Fabric Fails

Use this topic to learn about the troubleshooting scenario when a device addition to fabric fails.

Switch Addition to Fabric Fails

The following symptoms were observed when switch addition to fabric fails:

- SuperSpine close error
- SuperSpine Device x.x.x.x not connected to Spine Device y.y.y.y

```
SuperSpine clos error SuperSpine Device x.x.x.x not connected to Spine Device y.y.y.y

```

The issue occurred because there is a connection issue between the spine and super-spine device.

Resolution

Resolve the connection issue between the spine and super-spine devices.



Note

This error only occurs when attempting to add the spine device while there is a connection issue.

Addition of Multiple Devices to Fabric Fails

Adding multiple devices to fabric results in failure due to an OAuth2 PKI certificate issue.

```
efa fabric device add-bulk failed because of OAuth2 PKI certificate failed

```

The log entry indicates issue with incorrect username or password:

```
%!Error: Importing OAuth2 PKI certificate failed. Please verify certificate location and user credentials/parameters.

```

Resolution

Enter the correct username and password.

EFA Status Shows Nothing and gofaultmanager and gopolicy POD Services Stuck in init State

In a multi-node deployment, the following symptoms are observed:

1. The **efa stauts** command does not show any output.

```
extreme@tin10tpvm01:~$ efa status
+-----+-----+-----+-----+
| Node Name | Role | Status | IP |
+-----+-----+-----+-----+
--- Time Elapsed: 10.468338371s ---
```

2. Readiness failed for kube-system (POD calico-kube-controllers-685b65ddf9-npsxj is Unhealthy).
3. The pods gopolicy-service-v5wfz and gofaultmanager-service-vr7db are stuck in the initialization phase.

```
k3s kubectl -n efa get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP
NODE    NOMINATED NODE   READINESS GATES
gopolicy-service-v5wfz              0/1     Init:0/3    0         20d   10.42.194.102
efa      <none>          <none>
gosystem-service-mkvwm             1/1     Running    0         20d   10.42.194.96
efa      <none>          <none>
gosnmp-service-blwlx               1/1     Running    0         20d   192.168.246.10
efa      <none>          <none>
efa-api-docs-65bdp                 1/1     Running    0         20d   10.42.194.98
efa      <none>          <none>
gofaultmanager-service-vr7db       0/1     Init:0/2    0         20d   10.42.194.118
efa      <none>          <none>
```

The following reasons are observed for the failure:

- The following message appears for StorageUtilizationFullAlert under the gofaultmanager-service folder:

```
faultmanager/efa_gofaultmanager-service-9175g_db61b2b9-4e4e-4acb-b71d-e89ec6fcd0e8/
gofaultmanager-service/5.log
2023-08-15T07:08:06.751079487-05:00 stdout F
@time="2023-08-15T07:08:06.743764 CDT" level=info msg="Element
is :domain.AlertInventory{Name:\"StorageUtilizationThresholdAlert\", AlertID:0x7940,
Severity:\"Warning\", Resource:\"/App/System/Storage\", DisplayName:\"Storage
utilization threshold alert\"}"
2023-08-15T07:08:06.751104919-05:00 stdout F @time="2023-08-15T07:08:06.744859 CDT"
level=info msg="Element is :domain.AlertInventory{Name:\"StorageUtilizationFullAlert
\", AlertID:0x7941, Severity:\"Critical\", Resource:\"/App/System/Storage\",
DisplayName:\"Storage utilization full alert\"}"
```

- To verify, check the output of the `df -h` command on both nodes.

Resolution

Complete the following steps to recover the failure:

1. Failover to standby to recover EFA or XCO if disk space is only full on the active node.

```
sudo reboot -f now
```

2. After recovery, recover the disk space.
 - Run the `df -h` command on both nodes.
 - Run the `efactl clean` command to free disk space.

- Run the `df -h` command on both nodes again to confirm space availability.
3. Verify if the pods are stuck in evicted state, or showing `ImagePullBackOff` / `ImagePullNever`.

For more information, see [EFA or XCO K3s Pods Stuck in Evicted State](#).

4. Restart the EFA and k3s services.

```
efactl restart
systemctl restart k3s
```

5. If none of the above works, reinstall EFA or XCO via SLX-OS.

```
no efa deploy
efa deploy
```



Note

When a system is low on disk space, k3s handles the disk pressure event in two steps:

1. **Eviction of scheduled pods:** If disk space is cleaned up, the pods will be rescheduled by k3s.
2. **Image Deletion:** If disk space remains low after pod eviction, the images are marked for deletion. After cleanup, without a registry to import from, image pull fails when pods try to restart.

EFA Commands Fail to Execute

Running EFA commands fails with the following symptoms:

- EFA commands consistently fail to execute, resulting in errors.
- Switches become stuck in a locked state within EFA
- Multiple PODs enter Error and CrashLoopBackOff states

The issue occurred because the K3s CoreDNS POD experienced a critical error, disrupting communication between RabbitMQ and other PODs, leading to widespread restarts of PODs and system instability.

Resolution

To restore the system functionality, perform a failover from the active to the standby node.



Note

- Regular system monitoring is crucial for prompt issue detection and resolution.
- Engaging Global Technical Assistance Center (GTAC) support in a timely manner is essential for effective Root Cause Analysis (RCA) and issue resolution.

EFA Commands Fails with "Error 408"

The **efa inventory device setting update** command fails with a "408 Error" response.

The command cannot be completed at this time. Please check the node status using 'efa status' and retry the command later.

The issue is caused by an intermittent timeout of the efa-client connection request to the EFA or XCO node VIP on HTTP port 80.

Resolution

Upgrade to XCO 3.5.0 and later.

EFA Command Fails with "Dial TCP xx.xx.xx.xx:80: Connect: Connection Refused" Message

The different EFA commands fail with the following message:

```
'efa tenant create --name vpod01 --port 10.10.10.10[0/1:3,0/1:1] --vlan-range 200 --
enable-bd=true'
Command Response:
Error: Post "http://10.10.10.10/v1/tenant/tenant": dial tcp : connect: connection refused
```

This can also be preceded by the response `Error: Bad Gateway`.

The error occurs because the EFA CLI sends a REST POST message to the tenant (for `efa tenant create`), and the tenant is not responding, possibly because it is not running.

Resolution

Ensure that the tenant is running and responding to the POST messages.

EFA Error When Registering a RELP Handler on Endpoint Using FQDN

Adding subscriber for RELP with FQDN fails resulting in `The host from the endpoint URL is not reachable error`.

```
(efa:extreme)extreme@tpvm93:~$ efa notification subscribers add-syslog-relp --address
win01.etsuklab.com --insecure
Error: Error on registration for a 'relp' handler on endpoint 'win01.etsuklab.com:514'.
ERROR: The host from the endpoint URL is not reachable.
```

The issue occurred because `gonotification pod` is unable to resolve the FQDN to an IP address due to lack of DNS server entry in `/etc/resolv.conf` at the time of EFA or XCO installation.

Resolution

Upgrade to XCO 3.2.0 or above which allow coredns to pass the /etc/resolv.conf configuration.

```
(efa:extreme)extreme@tpvm203:/apps/efa$ k3s kubectl -n kube-system describe cm coredns |
grep resolv
#forward . /etc/resolv.conf
(efa:extreme)extreme@tpvm203:/apps/efa$ sudo ./update-dns.sh --dns-action allow
DNS entries forwarded to the pods now allowed
(efa:extreme)extreme@tpvm203:/apps/efa$ k3s kubectl -n kube-system describe cm coredns |
grep resolv
forward . /etc/resolv.conf
```

EFA Fabric Show Command Fails After Active Node Restart Causing CFG Refreshed Error on Devices

After restarting MariaDB on the active XCO node, the **efa fabric show** command returns an error message indicating a refused connection and invalid transaction

```
Error : dial tcp <xco-ip>:3306: connect: connection refused; invalid transaction; invalid transaction
```

Once the database connection is reestablished, one of the border leaf devices shows a `cfg-refresh-error`.

```
(efa:extreme)extreme@slx-tpvm:$ efa fabric show

| 192.168.246.21 |      | SPINE01      | 64512 | Spine      | provisioned | cfg refreshed |
LA,LD,IU      | BGP-D,INTIP-C,INTIP-D | NA | 1 |
| 192.168.246.22 |      | SPINE02      | 64512 | Spine      | provisioned | cfg refreshed |
LA,LD,IU,BGPU | BGP-D,INTIP-C,INTIP-U,INTIP-D | NA | 1 |
| 192.168.246.31 |      | LEAF001      | 65000 | Leaf       | provisioned | cfg in-sync   |
NA            | NA            | 2 | 1 |
| 192.168.246.32 |      | LEAF002      | 65000 | Leaf       | provisioned | cfg in-sync   |
NA            | NA            | 2 | 1 |
| 192.168.246.33 |      | LEAF003      | 65001 | Leaf       | provisioned | cfg in-sync   |
NA            | NA            | 2 | 1 |
| 192.168.246.34 |      | LEAF004      | 65001 | Leaf       | provisioned | cfg refreshed |
LA,LD,IU,POU  | MCT-U,MCT-PA,BGP-D,INTIP-C,INTIP-U,INTIP-D | 2 | 1 |
| 192.168.246.35 |      | BLEAF01      | 66000 | BorderLeaf | provisioned | cfg refresh error |
LA,LD,IU,POU  | BGP-D,INTIP-C,INTIP-D | 2 | 1 |
| 192.168.246.36 |      | BLEAF02      | 66000 | BorderLeaf | provisioned | cfg refreshed |
LA,LD,IU      | BGP-D,INTIP-C,INTIP-D | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
```

The issue occurred because the Inventory database was not fully updated with the device status after restarting MariaDB.

Resolution

To correct the app-state for the affected device, complete the following steps:

1. Shut down the MCT ports on SLX.
2. Update the device inventory on XCO by running the following command:

```
efa inventory device update -ip <device-ip>
```

3. Restart the MCT ports on SLX.

- Update the device inventory on XCO again by running the following command:

```
efa inventory device update -ip <device-ip>
```

- If the issue persists, perform a Drift Reconcile on the affected device using the following command:

```
efa inventory drift-reconcile execute --ip <device-ip> --reconcile
```

EFA or XCO is Down on Standby TPVM

Running the **efa status** command shows that the EFA or XCO is down on standby TPVM node.

```
extreme@tpvm01:~$ efa status
-----+
Node Name   Role      Status    IP
-----+
tpvm02     active    up        x.x.x.x
-----+
tpvm01     standby down    x.x.x.y
-----+
```

MariaDB service status: The the service is *loaded* and *enabled* but is currently *failed* with an exit code, indicating that the MariaDB service encountered an error and stopped running at a specific time

```
extreme@tpvm01:~$ systemctl status mariadb.service
● mariadb.service - MariaDB 10.6.9 database server
   Loaded: loaded (/lib/systemd/system/mariadb.service; enabled; vendor preset: enabled)
   Drop-In: /etc/systemd/system/mariadb.service.d
            └─migrated-from-my.cnf-settings.conf
   Active: failed (Result: exit-code) since Wed 2024-05-22 14:51:15 GMT; 55s ago
     Docs: man:mariabdb(8)
           https://mariadb.com/kb/en/library/systemd/
   Process: 30593 ExecStartPre=/bin/sh -c [ ! -e /usr/bin/galera_recovery ] && VAR= ||
   VAR=`cd /usr/bin/..; /usr/bin/galera_recovery`; [ $? -eq 0 ]   && systemctl set-env
   Process: 30591 ExecStartPre=/bin/sh -c systemctl unset-environment
   _WSREP_START_POSITION (code=exited, status=0/SUCCESS)
   Process: 30583 ExecStartPre=/usr/bin/install -m 755 -o mysql -g root -d /var/run/mysqlqld
   (code=exited, status=0/SUCCESS)
```

Journalctl output: Errors and notes indicate InnoDB plugin initialization failure, corrupted pages, and storage engine failure .

```
extreme@tpvm01:~$ journalctl -xe
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] InnoDB: Set
innodb_force_recovery=1 to ignore corruption.
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] InnoDB: Failed to read page
91 from file './dcapp_raslog/syslog.ib
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [Note] InnoDB: Set
innodb_force_recovery=1 to ignore corrupted pages.
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] InnoDB: Plugin
initialization aborted with error Generic error
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [Note] InnoDB: Starting shutdown...
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] Plugin 'InnoDB' init
function returned error.
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] Plugin 'InnoDB'
registration as a STORAGE ENGINE failed.
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [Note] Plugin 'FEEDBACK' is
disabled.
```

```

May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] Unknown/unsupported storage
engine: innodb
May 22 14:53:04 tpvm01 sh[508]: 2024-05-22 14:53:04 0 [ERROR] Aborting'
May 22 14:53:04 tpvm01 systemd[1]: mariadb.service: Control process exited, code=exited
status=1
May 22 14:53:04 tpvm01 systemd[1]: mariadb.service: Failed with result 'exit-code'.
May 22 14:53:04 tpvm01 systemd[1]: Failed to start MariaDB 10.6.9 database server.
-- Subject: Unit mariadb.service has failed
-- Defined-By: systemd
-- Support: http://www.ubuntu.com/support
--
-- Unit mariadb.service has failed.
--
-- The result is RESULT.
May 22 14:53:05 tpvm01 Keepalived_vrrp[2863]: Script `chk_mariadb` now returning 1
May 22 14:53:05 tpvm01 Keepalived_vrrp[2863]: VRRP_Script(chk_mariadb) failed (exited
with status 1)
May 22 14:53:05 tpvm01 Keepalived_vrrp[2863]: (HA1) Changing effective priority from 104
to 102

./efa_2024-05-21T14-10-31.983/efactl/efactl_x.x.x.y_20240521_debug.log
} , {
  "name": "tpvm01",
  "status": "down",
  "IP": "x.x.x.x",
  "role": "standby",
  "galera_size": "",
  "galera_host": false,
  "reason": "Host is not a member of Galera cluster"
} ] }'

```

The issue occurs because MariaDB is down on the standby TPVM.

Resolution

1. Restart MariaDB.

```
sudo systemctl restart mariadb
```

2. Restart the EFA or XCO service.

```
efactl restart
```

3. Restart TPVM.

```
tpvm stop
tpvm start
```

4. If none of the above steps work, consider reinstalling EFA or XCO.



Note

The following are some additional troubleshooting steps:

1. Stop MariaDB service.

```
sudo systemctl stop mariadb.service
```

2. Remove the Galera Cluster Cache.

```
sudo rm -f /apps/efadata/mysql/galera.cache
```

3. Restart MariaDB service.

```
sudo systemctl start mariadb.service
```

4. Confirm MariaDB has started.

```
systemctl status mariadb
```

EFA Fabric health Not in Sync with EFA Fabric Topology Show

The **efa fabric health** reports health status as Black for boarder leaves and is not in sync with the **efa fabric topology show**.



Note

In the following example, the Black status was due to the MCT ICL Port-channel64 being set to admin down. Although it was later brought admin up, the health status remained Black.

```
Fabric Name           : dcb11_cnis
Fabric Type           : clos
Fabric Health         : Black
Fabric Status         : configure-success
Fabric Level Physical Topology Health : Green
Fabric Device Health
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | ROLE | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.152.68.112 | Leaf | Black | Black | Black |
| 10.152.68.109 | Leaf | Green | Green | Green |
| 10.152.68.114 | Spine | Black | Black | Black |
| 10.152.68.110 | Leaf | Green | Green | Green |
| 10.152.68.113 | Spine | Black | Black | Black |
| 10.152.68.111 | Leaf | Black | Black | Black |
| 10.152.68.122 | BorderLeaf | Black | Black | Black |
| 10.152.68.121 | BorderLeaf | Black | Black | Black |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

In the following **efa fabric topology show** command output, the BGP neighbor relationship between the border leaves over Port-channel64 is UP and ESTABLISHED.

```
(efa:extreme)extreme@tpvm941:~$ efa fabric topology show underlay --name dcb11_cnis
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP | SOURCE DEVICE ASN |
| DESTINATION DEVICE ASN | NEIGHBOR AFI STATE | NEIGHBOR SAFI | UNDERLAY STATE |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.152.68.122 | 10.152.68.114 | 172.31.254.107 | 10.10.10.10 | 4200003000 | |
| 4200002000 | | ipv4 | unicast | ESTAB | |
| 10.152.68.122 | 10.152.68.114 | 172.31.254.107 | 10.10.10.10 | 4200003000 |
| 4200002000 | | l2vpn | evpn | ESTAB | |
| 10.152.68.122 | 10.152.68.113 | 172.31.254.107 | 10.10.10.4 | 4200003000 |
| 4200002000 | | ipv4 | unicast | ESTAB | |
| 10.152.68.122 | 10.152.68.113 | 172.31.254.107 | 10.10.10.4 | 4200003000 |
| 4200002000 | | l2vpn | evpn | ESTAB | |
| 10.152.68.122 | 10.152.68.113 | 172.31.254.107 | 10.10.10.6 | 4200003000 |
| 4200002000 | | ipv4 | unicast | ESTAB | |
| 10.152.68.122 | 10.152.68.113 | 172.31.254.107 | 10.10.10.6 | 4200003000 |
| 4200002000 | | l2vpn | evpn | ESTAB | |
| 10.152.68.122 | 10.152.68.121 | 172.31.254.107 | 10.20.20.2 | 4200003000 |
| 4200003000 | | ipv4 | unicast | ESTAB | |
| 10.152.68.122 | 10.152.68.114 | 172.31.254.107 | 10.10.10.8 | 4200003000 |
| 4200002000 | | ipv4 | unicast | ESTAB | |
| 10.152.68.122 | 10.152.68.114 | 172.31.254.107 | 10.10.10.8 | 4200003000 |
| 4200002000 | | l2vpn | evpn | ESTAB | |
| 10.152.68.121 | 10.152.68.114 | 172.31.254.180 | 10.10.10.12 | 4200003000 |
| 4200002000 | | ipv4 | unicast | ESTAB | |
| 10.152.68.121 | 10.152.68.114 | 172.31.254.180 | 10.10.10.12 | 4200003000 |

```

```

| 4200002000          | l2vpn          | evpn          | ESTAB          |
| 10.152.68.121     | 10.152.68.113 | 172.31.254.180 | 10.10.10.3    | 4200003000
| 4200002000          | ipv4           | unicast       | ESTAB          |
| 10.152.68.121     | 10.152.68.113 | 172.31.254.180 | 10.10.10.3    | 4200003000
| 4200002000          | l2vpn          | evpn          | ESTAB          |
| 10.152.68.121     | 10.152.68.113 | 172.31.254.180 | 10.10.10.1    | 4200003000
| 4200002000          | ipv4           | unicast       | ESTAB          |
| 10.152.68.121     | 10.152.68.113 | 172.31.254.180 | 10.10.10.1    | 4200003000
| 4200002000          | l2vpn          | evpn          | ESTAB          |
| 10.152.68.121     | 10.152.68.122 | 172.31.254.180 | 10.20.20.3    | 4200003000
| 4200003000          | ipv4           | unicast       | ESTAB          |
| 10.152.68.121     | 10.152.68.114 | 172.31.254.180 | 10.10.10.14   | 4200003000
| 4200002000          | ipv4           | unicast       | ESTAB          |
| 10.152.68.121     | 10.152.68.114 | 172.31.254.180 | 10.10.10.14   | 4200003000
| 4200002000          | l2vpn          | evpn          | ESTAB          |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+

```

The health state for devices reports ‘Black’ because the connection states in the inventory database have not been refreshed, whereas ‘fabric topology show’ fetches connection details from devices dynamically, showing ‘ESTABLISHED’.

Resolution

Upgrade to XCO 3.3.1 or above. In XCO 3.3.1, the **fabric topology show** command fetches connection details from the inventory database, eliminating any discrepancy between **efa fabric topology show** and **efa fabric health**.

EFA or XCO VM is down and Came Up After Stopping the Services

EFA or XCO VM is down, but it came up after stopping the services. The following symptoms are observed:

1. Only two out of three PODs are running; the third is in the Init state.
2. Not all PODs are available.

```

(efa)efaadmin@$ sudo efactl status
Node: efa
NAME READY STATUS RESTARTS AGE IP NODE NOMINATED NODE READINESS GATES
pod/efa-api-docs-qgdf8 1/1 Running 0 2m31s 10.42.194.81 efa <none> <none>
pod/rabbitmq-7pkhz 1/1 Running 0 2m31s 10.42.194.83 efa <none> <none>
pod/gopolicy-service-29p4s 0/1 Init:2/3 0 2m31s 10.42.194.82 efa <none> <none>

```

It is further observed that not all the pods booted up after the upgrade.

Resolution

No issues observed after reboots.

EFA or XCO Execute-CLI Reports SLX Switches as Unreachable

The following output is an example of a failed execute-cli with devices not reachable:

```

(efa:extreme)extreme@:~$ efa inventory device execute-cli --ip 1.1.1.1 --command "show clock"
Execute CLI [failed]
+-----+-----+-----+-----+-----+-----+-----+-----+
| IP Address | Host Name | Fabric | Command | Status | Reason | Output |

```

```

+-----+-----+-----+-----+-----+-----+
| 1.1.1.1 | test   |      | show clock | Failed | Device 1.1.1.1 not reachable. |      |
|         |       |     |           |       | Please retry after verifying the |     |
|         |       |     |           |       | inputs and connectivity issues. |     |
+-----+-----+-----+-----+-----+

```

This issue occurs due to a password change. The old password is being used despite the change.

Resolution

Use the correct password after the password change.

EFA or XCO Execute-CLI Fails on SLX

On the SLX platform, EFA or XCO is unable to retrieve command output from the SLX device. The error log is as follows:

```

{"@time":"2022-11-07T09:16:01.829058 PST","App":"inventory","level":"error","msg":"Failed
to establish connection with device(10.247.33.21), Error: Device 10.247.33.21 not
reachable.
Please retry after verifying the inputs and connectivity
issues.", "rqId":"377e8e2a-83f7-4c09-9fbb-50fa3298e5a9"}

```

It was found that commands are executed in parallel mode rather than serial mode in the following sequence:

Script tool (Python script) > EFA > SLX

Resolution

When using parallel connections, HTTP requests are sent simultaneously without waiting for replies. This causes commands to get queued up and may lead to random execution failures.

Removing EFA or XCO from the path enables commands to run in serial order, which resolves the issue.

EFA or XCO Fails to Resolve Non-Cluster FQDN

The EFA or XCO is unable to resolve the FQDN of an external server. When running the following command:

```

efa:extreme)extreme@tpvmt:~$ k3s kubectl exec pods/goauth-service-68ff9 -n efa --
nslookup google.com

```

The output shows the following error:

```

Defaulted container "go-auth" out of: go-auth, check-db-ready (init), wait-for-rabbitmq
(init)
Server:      10.43.0.10
Address:     10.43.0.10:53

** server can't find google.com: SERVFAIL

** server can't find google.com: SERVFAIL

```

```
command terminated with exit code 1
```

The issue occurs because the CoreDNS lacks the necessary DNS entry.

Resolution

Add the DNS entry in the `/etc/resolv.conf` file and upgrade to EFA 3.0 or later. The forward entry will automatically be added to CoreDNS.



Note

To inspect the CoreDNS configuration, run the following command:

```
extreme@tpvmt:~$ k3s kubectl describe cm coredns -n kube-system
```

The following is an example output details:

```
Name:          coredns
Namespace:     kube-system
Labels:        <none>
Annotations:   <none>

Data
====
Corefile:
-----
.:53 {
  errors
  health
  ready
  kubernetes cluster.local in-addr.arpa ip6.arpa {
    pods insecure
    fallthrough in-addr.arpa ip6.arpa
  }
  hosts /etc/coredns/NodeHosts {
    ttl 60
    reload 15s
    fallthrough
  }
  prometheus :9153
  forward . /etc/resolv.conf
}
```

EFA or XCO Fails to Log in to Switch

Unable to access EFA or XCO on switch due to an invalid password.

Resolution

Reset the password.

Fabric Skill Troubleshooting

Issue	Resolution
<p>Physical topology is wrong. They are wrongly cabled. How do I correct the physical topology?</p>	<ol style="list-style-type: none"> 1. Correct the Physical Topology 2. Update the inventory using the efa inventory device update command. 3. Configure the fabric using the efa fabric configure command. 4. If the above operations fail, perform the following steps for the newly added device: <ol style="list-style-type: none"> a. Remove the newly added device from the fabric using the efa fabric device remove --ip <device-ip> --name <fabric-name> command. b. Ensure that the Physical Topology is correct. c. Update the Inventory using the efa inventory device update command. d. Add the device back to the fabric using the efa fabric device add command. e. Reconfigure the fabric using the ea fabric configure command.
<p>Unable to configure a Non-Clos fabric whose devices in a rack are not connected to each other.</p>	<ul style="list-style-type: none"> • Non-Clos fabric does not allow single-homed devices in a rack. Before you add these devices, ensure that you have a connection between them. • If you already have such devices in a fabric then remove the devices, make the connection, update the inventory, and add them back to the fabric.

Issue	Resolution
<p>Devices are in "cfg-refreshed" state</p>	<ul style="list-style-type: none"> • If you have not triggered the manual DRC, the efa fabric show --name <fabric name> command shows devices in "cfg-refreshed" state, and the intended BFD configuration does not get configured on the devices. • This is due to config mismatch between the device and the expected fabric config, this usually gets auto corrected whenever the rediscovery of devices happen in the backend; if it does not it can be manually reconciled with the efa fabric debug device drift --name <fabric-name> --deviceip <device-ip> --reconcile command. • Before the devices are reconciled, identify the drift by running the efa fabric debug device drift --name <fabric-name> --device-ip <deviceip> command. <p>Search the debug logs with the string "DriftCompare" to get the fabric service and asset service params comparison details that identified the drift. For example,</p> <pre># cat /var/log/efa/fabric/fabric-server.log grep "DriftCompare" ... {"@time":"2022-05-05T04:00:08.711548 PDT","level":"info","msg":"DriftCompare: compareInterfaceConfigs Refreshed command SET IP:64:portchannel: 10.20.20.2/31 , portChannelIPMissing asset, assetIP : "}</pre>
<p>Unable to add or configure fabric due to existing configs on the device</p>	<ul style="list-style-type: none"> • Before you configure the fabric, ensure that the devices are cleaned. • If the fabric configuration fails, remove the device from fabric using the efa fabric device remove command. • Run the clear config command on device using the efa fabric debug clear-config command. • Update the devices on inventory. • Add devices and configure fabric. • If any device in a fabric is in "admin-down" state, the following commands in the same fabric will not add or delete devices: efa fabric device add-bulk and efa fabric device remove.
<p>Unable to configure a Clos fabric whose leaf or border-leaf devices are not connected to Spine.</p>	<ul style="list-style-type: none"> • Remove devices from the fabric. • Ensure that the LLDP config on leaf or border-leaf devices is enabled for Spine. • Update all the devices in inventory. • Add or configure the Clos fabric.

Issue	Resolution
<p>When a fabric device is in admin-down state, removing it from inventory does not clear its configurations.</p>	<ul style="list-style-type: none"> • This is expected behavior. • If any device in a fabric is in "admin-down" state, the efa fabric device add-bulk and efa fabric device remove commands in the same fabric do not add or delete devices. • Wait for the device to admin-up. Check the device status in inventory using the efa inventory device list command. • Delete device from inventory.
<p>CLI is blocked from processing any fabric commands.</p>	<p>Run the efa fabric debug service lock command.</p> <pre> efa fabric debug service lock +-----+-----+-----+ Lock type Locked Reason +-----+-----+-----+ REST Lock false +-----+-----+-----+ Service Lock false +-----+-----+-----+ DB Lock false +-----+-----+-----+ Device 10.20.244.200 false +-----+-----+-----+ Device 10.20.244.201 false +-----+-----+-----+ Lock Status </pre>
<p>How to check if an event is received at fabric service.</p>	<p>Check the following string in the fabric debug logs:</p> <pre> # cat /var/log/efa/fabric/fabric-server.log grep -i Recei grep -i message {"@time":"2022-05-05T03:59:59.557814 PDT","level":"info","msg":"Received BlockedProcessEventMsg message"} </pre>

Fabric Skill is Down After a Failed Token Certificate Renewal

Versions prior to EFA or XCO 3.x.x did not include the "kustomization.yaml" file.

After upgrading a system from EFA 2.x.x to EFA or XCO 3.x.x, attempts to update or refresh the token certificate fails due to the missing "kustomization.yaml" file in the directories:

- **Server**

`/opt/efa/certs/cert`

`/opt/efa/certs/key`

- **TPVM**

`/apps/efa/certs/cert`

`/apps/efa/certs/key`

This results in EFA or XCO entering a down state.



Note

This issue does not occur in XCO 3.5.0 or above.

The "kustomization.yaml" file is incorrectly created in a nested folder:

- /[opt|apps]/efa/certs/certs/cert/
- /[opt|apps]/efa/certs/certs/key/

This issue does not affect the EFA upgrade process or existing certificates but is only encountered when attempting to renew the Token Certificate.

When attempting to refresh the certificate type "token" in XCO, the renewal fails and commands will no longer process:

```
(efa:administrator)administrator@GTAC-SP-EFA-1:~$ efa certificate server renew --cert-type token
Error: Certificate renewal has failed.

Renew certificates for EFA

Usage:
  efa certificate server renew [flags]

Flags:
  --cert-type string    Type of certificate to renew, valid values (server | token | root-ca | intermediate-ca | k3s-ca | k3s-server ) (default "server")
  --- Time Elapsed: 2m47.386569066s ---
```

GoAuth Pod error messages:

```
efa      25m      Warning   FailedMount   pod/goauth-service-bs7j2   MountVolume.SetUp failed for volume "verifier" : configmap "jwt-verifier" not found
efa      25m      Warning   FailedMount   pod/goauth-service-bs7j2   MountVolume.SetUp failed for volume "signer" : secret "jwt-signer" not found
efa      25m      Warning   FailedMount   pod/goauth-service-bs7j2   Unable to attach or mount volumes: unmounted volumes=[verifier signer], unattached volumes=[], failed to process volumes=[]: timed out waiting for the condition
```

Installer logs:

```
+(2024-02-15T21:32:36.304475 UTC GTAC-SP-EFA-1 common.sh:3130):
token_keypair_regenerate(): echo 'Generating yaml files for certificates.'
Generating yaml files for certificates.
+(2024-02-15T21:32:36.305336 UTC GTAC-SP-EFA-1 common.sh:3131):
token_keypair_regenerate(): reliable_k3s kubectl kustomize /opt/efa/certs/cert
+(2024-02-15T21:32:36.306194 UTC GTAC-SP-EFA-1 common.sh:20845): reliable_k3s():
without_debug reliable_k3s_impl kubectl kustomize /opt/efa/certs/cert
error: unable to find one of 'kustomization.yaml', 'kustomization.yml' or 'Kustomization' in directory '/opt/efa/certs/cert'
+(2024-02-15T21:32:36.379913 UTC GTAC-SP-EFA-1 common.sh:3132):
token_keypair_regenerate(): reliable_k3s kubectl kustomize /opt/efa/certs/key
+(2024-02-15T21:32:36.380757 UTC GTAC-SP-EFA-1 common.sh:20845): reliable_k3s():
without_debug reliable_k3s_impl kubectl kustomize /opt/efa/certs/key
error: unable to find one of 'kustomization.yaml', 'kustomization.yml' or 'Kustomization' in directory '/opt/efa/certs/key'
+(2024-02-15T21:32:36.455284 UTC GTAC-SP-EFA-1 common.sh:3134):
token_keypair_regenerate(): reliable_k3s kubectl apply -f /opt/efa/certs/key/cert-secret.yaml -n efa
```

```
+ (2024-02-15T21:32:36.456091 UTC GTAC-SP-EFA-1 common.sh:20845): reliable_k3s():  
without_debug reliable_k3s_impl kubectl apply -f /opt/efa/certs/key/cert-secret.yaml -n  
efa  
error: no objects passed to apply  
+ (2024-02-15T21:32:36.623988 UTC GTAC-SP-EFA-1 common.sh:3135):  
token_keypair_regenerate(): reliable_k3s kubectl apply -f /opt/efa/certs/cert/cert-  
configmap.yaml -n efa  
+ (2024-02-15T21:32:36.624898 UTC GTAC-SP-EFA-1 common.sh:20845): reliable_k3s():  
without_debug reliable_k3s_impl kubectl apply -f /opt/efa/certs/cert/cert-configmap.yaml  
-n efa  
error: no objects passed to apply
```

Workaround

1. Complete any of the following:

- Option 1:** Move the "kustomization.yaml" files to the correct locations as shown above.
- Option 2:** If the files do not exist, create them manually:



Note

- Ensure the name is the same for both files, but their contents are unique for the cert/ and key/ files.
- Use the server path /opt/efa/. If using TPVM, replace it with /apps/efa/.

i. Create the kustomization.yaml file for certificate:

```
administrator@GTAC-SP-EFA-1:~$ sudo vi /opt/efa/certs/cert/kustomization.yaml  
  
configMapGenerator:  
- name: jwt-verifier  
  files:  
  - cert.crt  
generatorOptions:  
  disableNameSuffixHash: true
```

ii. Create the kustomization.yaml file for key:

```
administrator@GTAC-SP-EFA-1:~$ sudo vi /opt/efa/certs/key/kustomization.yaml  
  
secretGenerator:  
- name: jwt-signer  
  files:  
  - sign.key  
generatorOptions:  
  disableNameSuffixHash: true
```

iii. After restoring the files, if XCO is in a trouble state, continue with the following:

If the files are corrected proactively and XCO is stable, no need to continue.

2. Run the **efa_renew_certs.sh** script to complete the token renewal process.

```
administrator@GTAC-SP-EFA-1:/opt/efa$ sudo ./efa_renew_certs.sh --type token
```

Post token renewal process, GoAuth re-initializes, starts running, and then XCO returns to service (up and running).

Fabric Configuration Fails with EFA or XCO

SLX fabric fails to form when created through EFA or XCO using the following commands for a two-switch small data center fabric:

```
efa fabric create --name FabricName --type non-clos --description Row_A
efa fabric device add --ip 10.254.3.20 --hostname SLX_1 --rack A1 --username admin --password password --name FabricName
efa fabric device add --ip 10.254.3.30 --hostname SLX_2 --rack A2 --username admin --password password --name FabricName
```

The `efa fabric show` command output shows the following error:

```
$ efa fabric show --name FabricName

Fabric Name: FabricName, Fabric Description: Row_A, Fabric Type: non-clos
+-----+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | RACK | HOST NAME | ASN | ROLE | DEVICE STATE |
| APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.254.3.30 | A2 | SLX_2 | 4200000001 | leaf | not provisioned |
| cfg ready | DA | | SYSP-C,BGP-C,INTIP-C,EVPN-C,O-C | 2 | 1 |
| 10.254.3.20 | A1 | SLX_1 | 4200000000 | leaf | not provisioned |
| cfg ready | DD,DA | | SYSP-C,BGP-C,INTIP-C,EVPN-C,O-C | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+

efa fabric configure --name FabricName
Validate Fabric [Failed]
Missing Links
Rack A1, should have 2 number of devices
Rack A2, should have 2 number of devices
Configure Fabric [Failed]
Error: fabric validation failed
```

The described topology functions as a single ToR with redundancy, even though the devices are in different physical racks. As a best practice, ensure that there are two devices in each rack in a non-Clos fabric for redundancy.

Resolution

Configure the fabric with the same rack name or use the add-bulk EFA command to form the fabric.

- The following example configures a fabric using the same rack name:

```
efa fabric device add --ip 10.254.3.20 --hostname SLX_1 --rack A1 --username admin --password password --name FabricName
efa fabric device add --ip 10.254.3.30 --hostname SLX_2 --rack A1 --username admin --password password --name FabricName
```

- The following example configures a fabric using the add-bulk command:

```
efa fabric device add-bulk --ip 10.254.3.20,10.254.3.30 --username admin --password password --name FabricName --rack A1
```

How to Add a Second Link to an Existing Fabric via EFA or XCO

Complete the following steps to add a new link to an existing fabric using XCO or EFA.

1. Verify the fabric status.
 - Run the **efa fabric show** command to display the fabric status.
 - Run the **efa fabric error show --name <fabric_name>** command to check for errors in the existing fabric.
2. Review the existing physical and underlay topology.
 - Run the **efa fabric topology show physical --name <fabric_name>** command to display the physical topology.
 - Run the **efa fabric topology show underlay --name <fabric_name>** command to display the underlay topology.
3. Enable the new link.
 - Run the **efa inventory device interface set-admin-state --ip <IP> --if-type eth --if-name <interface> --state up** command to enable the new link.
 - Run the **efa fabric topology show underlay --name <fabric_name>** command to confirm the new link is added to the underlay topology.
4. Verify the link addition.
 - Run the **efa fabric topology show physical --name <fabric_name>** command to confirm the new link is added to the physical topology.
5. Configure the fabric.
 - Run the **efa fabric configure --name <fabric_name>** command to add the new link to the fabric.
6. Check the fabric status.
 - Run the **efa fabric show** command to verify the fabric status.
7. Verify the interface configuration.
 - Check the interface configuration on the device to ensure the new link is properly configured.



Note

Replace <fabric_name> with the actual name of your fabric, <IP> with the actual IP address of the device where the link is added, and <interface> with the actual interface name.

Log in to EFA or XCO Fails

The login to EFA or XCO fails. The supportsave data revealed the following entries in the database/*_error.log file:

```
2022-04-22 14:15:51 0 [Note] InnoDB: Buffer pool(s) load completed at 220422 14:15:51
2022-04-22 14:17:25 1 [ERROR] Slave SQL: Error 'Duplicate key name 'kine_name_index''
on query. Default database: 'k3s'. Query: 'create index kine_name_index on kine (name)',
Internal MariaDB error code: 1061
2022-04-22 14:17:25 1 [Warning] WSREP: Ignoring error 'Duplicate key name
```

```
'kine_name_index' on query. Default database: 'k3s'. Query: 'create index
kine_name_index on kine (name)', Error_code: 1061
```

The log file indicates a duplicate key name `kine_name_index` in the `k3s` database. This error suggests that the Kubernetes state stored in MariaDB is incorrect.

It is found that the Kubernetes state is stored in MariaDB and synchronized across both nodes of a high-availability (HA) cluster using Galera. The error indicates an inconsistency in the database state.

Resolution

Completely reset Kubernetes and then rebuild the Kubernetes configuration based on the rules used by the EFA or XCO installer.



Note

The business data, such as switch configurations and fabrics, are stored in separate database tables and are not modified by this procedure. However, consider taking an EFA or XCO backup as a precaution, if possible.

For further assistance with the recovery process, contact Extreme GTAC.

LDAP Authentication Configuration Using FQDN Fails

An attempt to log in to LDAP after configuring authentication with FQDN fails. Initially, configuration failed due to an outdated `/etc/resolv.conf` file. However, after running the `sudo <location of the script>/update-dns.sh --dns-action allow` command, LDAP authentication configuration using FQDN was successful. Now, login attempts fail with the following error log:

```
==> /apps/efa_logs/auth/auth-server_err.log <==
{"@time":"2023-09-15T13:14:38.029361 GMT","level":"error","msg":"Could not authenticate
against LDAP unable to read LDAP response packet: read tcp x.x.x.x:55630-
\u003ey.y.y.y:636: read: connection reset by peer"}
```

The issue occurred because of the port mismatch, as XCO uses port 389 for LDAP, whereas the configuration was set to port 636.

Resolution

Use the `--port 389` when logging in with the `efa login --username <>` command. A network connectivity test (`nc -zv y.y.y.y 389`) confirms that the connection to port 389 was established successfully.

```
extreme@<name-tpvm>:~$ nc -zv y.y.y.y 389
Connection to y.y.y.y 389 port [tcp/ldap] succeeded!
```

Management Subinterface IP Address Not Listed

In XCO 3.2.x, the main management interface's IP address is not displayed in the netstat output for port 443, unlike previous versions (EFA 3.1.x and earlier).

In EFA 3.1.x and earlier versions, the IP addresses are displayed as follows.

```
(efa:extreme)extreme@tpvmb:~$ netstat -utnlp | grep 443
(Not all processes could be identified, non-owned process info
will not be shown, you would have to be root to see it all.)
tcp        0      0 10.10.1.1:443      0.0.0.0:*           LISTEN      -
tcp        0      0 10.26.10.10:443    0.0.0.0:*           LISTEN      -
tcp6       0      0 :::6443            :::*                 LISTEN      -
```

In XCO 3.2.x, the IP addresses are not displayed as follows.

```
efa:extreme)extreme@tpvmt:~$ netstat -utnlp | grep 443
(Not all processes could be identified, non-owned process info
will not be shown, you would have to be root to see it all.)
tcp6       0      0 :::6443            :::*                 LISTEN
```

The behavior change is due to the upgrade to k3s latest version V1.25.6 in XCO 3.2.x.

Resolution

In older k3s versions, the TCP 443 service ran directly on the host, whereas in the current k3s version, the TCP 443 service runs within the traefik pod.

1. To verify, list pods using the **k3s kubectl get pods -n kube-system** command.
2. Verify the TCP 443 service on the Traefik pod using the **k3s kubectl -n kube-system exec <traefik-pod> -- netstat -tunlp | grep 443** command.

This must output `tcp 0 0 :::443 :::* LISTEN 1/traefik`, confirming the service is running in the traefik pod.

Recover a Leaf Node after Deletion

Deleting a leaf node using the **efa inventory device delete** command will remove the leaf node configuration.

Running the **efa inventory device delete --ip [leaf-node]** command removes the configuration applied by EFA or XCO on the leaf node. Exercise caution when using this delete option and know how to recover the unit if this action is performed accidentally. This procedure can also be useful in case of an RMA or if any configuration is lost upon an SLX reload.



Note

The **efa fabric show** and **efa inventory device list** commands will not display the leaf node after it has been removed with the delete command.

Resolution

1. Take a backup before deletion.

Before running the **efa inventory device delete --ip <ip address>** command, back up the current "good" state with the **efa system backup** command. If a manual backup is not done, use the most current automated system backup where the device was in a working state.

2. Restore EFA or XCO system from a backup.

Run the **efa system restore** command and select the desired backup.

3. Verify device status.

After the restore, the **efa device inventory list** and **efa fabric show** commands will report the device as added back to the inventory and fabric, but it will be in a `cfg-refresh-err` state.

4. Reconcile configuration.

- a. Run the **efa inventory drift-reconcile execute -ip <ip address> --reconcile** command for the device.
- b. Run the **efa inventory drift-reconcile history** command to see if the reconcile is completed.
- c. Run the **show running-config** command on the SLX-OS CLI to confirm that the configuration is applied on the leaf node.

5. Resolve `cfg-refresh-err`.

- If `cfg-refresh-err` is still shown for the spine devices and any MCT peer associated with the leaf device, perform the following:
 - Run the **efa inventory device update --fabric [fabric_name]** command.
 - If the `cfg-refresh-err` message persists, perform another reconcile on all impacted devices using the **efa inventory drift-reconcile execute -ip [IP_Address_comma_separated] --reconcile** command.

6. Verify that all the devices are in sync. Use the following commands to ensure all devices are in the `cfg in-sync` state.

- **efa fabric show**
- **efa inventory device list**
- **efa tenant po show**

SLX, EFA, or XCO Config Missing After SLX Reload

After reloading a leaf device, the cluster configuration and port-channel settings were lost.

The issue occurs in a non-standard management network setup, physically connecting their management interfaces to the cluster peer. This setup required the management network to traverse the MCT ICL. When a Drift Reconcile (DRC) is triggered, EFA or XCO shut down PO 64, which disrupted the management network and connectivity to the device being updated, resulting in an unexpected configuration state.

Resolution

EFA or XCO relies on the management network. If connectivity is lost to an IP fabric member during a DRC, the fabric can end up in an unexpected state. To ensure optimal EFA performance, use a dedicated out-of-band management network.

TPVM Unable to Authenticate to Red Hat Directory Service (LDAP)

An attempt to log in to the TPVM after configuring LDAP to a RHDS (Red Hat Directory Service) device, the RHDS server logs indicate incorrect credentials. However, after configuring the same server on XCO, the login is successful with the same credentials.

When authenticating to an LDAP server, multiple attributes are sent from the LDAP server to the TPVM to complete the login process. These include the user, password, home directory, and the desired shell environment. During an attempt to log in via "tpvm console," the following error message displayed after creating the home directory:

```
/bin/ksh: File not found
```

Resolution

Update RHDS options to send `/bin/bash` as the preferred shell.



Note

This value is cached and may take 5-10 minutes to update in the TPVM. Confirm the settings by logging in as "extreme" and by running the `sudo getent passwd [username]` command. For example,

```
extreme@tpvm21:~$ getent passwd extreme
extreme:x:1000:1000:extreme:/home/extreme:/bin/bash
```

Unable to Run EFA Commands

Unable to run EFA commands after the daylight savings time change. The following symptoms are observed:

1. Receiving `connection refused` error when running EFA commands.
2. Receiving `Unable to connect to the server: x509: certificate has expired or is not yet valid` error when running `k3s` command.

See [Unable to Connect to the Server](#) on page 71.

It is found that the time on the server does not match with `k3s` or `k8s`. This causes the certificate, used for validating calls, to be seen as invalid.

Resolution

Ensure that the EFA server is synchronized with an NTP server. If it's not configured, this error can occur.

To configure and sync the time to an NTP server, complete the following steps:

1. Configure and sync the time to an NTP server.

```
sudo apt-get update
sudo apt-get install ntp
sudo vi /etc/ntp.conf
```

2. Remove the pool lines and add the NTP server IP.

```
server 10.31.2.80
```

3. Restart and check the status of the NTP service.

```
sudo service ntp restart
sudo service ntp status
```

4. Set the timezone to UTC (Ubuntu has PDT timezone by default).

```
sudo timedatectl set-timezone UTC
```



Note

1. If the issue persists after time synchronization, collect the following information on the EFA server:

```
cd /var/lib/rancher/k3s/server/tls
```

2. Collect the following information for loop output:

```
cd /var/lib/rancher/k3s/server/tls
for i in `ls *.crt`; do
  echo $i;
  openssl x509 -startdate -noout -in $i;
done
date
```

This will help diagnose if the certificates are still seen as invalid due to time discrepancies.

Unable to Connect to the Server

When running the **k3s kubectl get pods -n efa** command to verify that all PODs are in a running state, the following error is displayed:

```
Unable to connect to the server: x509: certificate has expired or is not yet valid
```

```
(efa)extreme@slx-tpvm:~$ k3s kubectl get pods -n efa
Unable to connect to the server: x509: certificate has expired or is not yet valid:
current time 2020-06-10T13:59:48+02:00 is before 2023-06-09T09:12:24Z
```

It is found that either the certificate has expired or is not yet valid because the system date is incorrect.

```
(efa:extreme)extreme@SLX-tpvm:~$ date
Wed Jun 10 13:58:56 CEST 2020
```

Resolution

1. Renew the k3s-ca certificates which will also renew k3s-server certificates.

```
efa certificate server renew --cert-type k3s-ca
```

2. Renew the RootCA which will update intermediate and server certificates.

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

3. Renew JSON Web Token (JWT) Certificate:

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

4. Update device certificates for the registered devices:

```
efa certificate device install --ip x.x.x.x --force
efa inventory device update --ip x.x.x.x
```

Update Authentication Preference

By default, the authentication preference is set to 1 for HOST and 2 for LDAP. If both HOST and LDAP authentication are configured, HOST takes precedence, and LDAP only works when the HOST configuration is removed.

The system is configured for authentication with both local users (HOST) and LDAP users. HOST is the default first preference, and LDAP is secondary, which means LDAP will only work if the HOST user is removed. This is due to the default preference order of 1 for HOST and 2 for LDAP.

The default authentication preference prioritizes HOST. The order in which authentication methods are saved determines their preference. For example, if TACACS is saved first and then LDAP, the preference is set accordingly:

```
The default preference is HOST authentication.
The order is same in which user saves the other options.
In the example below, tacacs is saved first and then ldap is saved. The preference is set
in the same order:
(efa:extreme)extreme@tpvm-71:~$ efa auth authentication preference show
+-----+-----+-----+
| Auth Type | Identifier | Preference |
+-----+-----+-----+
| HOST      | HOST      | 1          |
+-----+-----+-----+
| TACACS    | 10.37.32.51 | 2          |
+-----+-----+-----+
| LDAP     | ldap2     | 3          |
+-----+-----+-----+
--- Time Elapsed: 123.236826ms ---

The user can update the preferences using this cli: efa auth authentication preference
update
```

Resolution

To update the authentication preference (or to change the default authentication preference from prioritizing HOST (1) over LDAP (2)), use the following commands:

- Add authentication preference.

```
efa auth authentication preference add --authType {TACACS | LDAP | LOCAL | HOST} --
identifier string --preference {1 | 2 | 3 | 4 | 5}
```

- Update authentication preference.

```
efa auth authentication preference update --authType {TACACS | LDAP | LOCAL | HOST}
--identifier string --preference {1 | 2 | 3 | 4 | 5}
```

- Delete authentication preference.

```
efa auth authentication preference delete --authType {TACACS | LDAP | LOCAL | HOST}
--identifier string --preference {1 | 2 | 3 | 4 | 5}
```

- Show the current authentication preference.

```
efa auth authentication preference show
```



Note

- For information about commands and supported parameters to configure authentication preference, see *ExtremeCloud Orchestrator Command Reference, 4.0.2*.
- For information about authentication policy CLI configuration, see *ExtremeCloud Orchestrator Security Configuration Guide, 4.0.2*.

The following is a sample procedure to set LDAP as the first preference and HOST as the second preference by running the following commands:

1.

```
efa auth authentication preference update --authType LDAP --identifier ldap2 --preference 1
```
2.

```
efa auth authentication preference update --authType HOST --identifier HOST --preference 2
```

This will change the authentication order, making LDAP the primary method and HOST the secondary method.

Update of Maximum Password Age of SLX Password Fails

The maximum password age for the SLX password is not updated on the SLX when using the following command:

```
efa inventory device secure settings update --max-password-age 0
efa inventory device secure settings apply --fabric BIS-CNIS-FABRIC
```

The issue occurred because the maximum password age value of 0 was previously set on XCO, but the values were reset to 90 on the SLX via the SLX-OS CLI.

Resolution

Run DRC with reconcile.

Analyze 10 Second Traffic Loss During DN Leaf Power Cycle

You experience a 10 second traffic loss at a node when the MCT peer was power-cycled to run a fail over.

Upon reviewing the logs, it appears that during the failover initiated by shutting down SLX_L01, BGP X.X.X.3 was affected, which in turn impacted the MCT peer of SLX_L02, X.X.X.2.

The following events occurred:

- ICL ports and cluster peer went down.
- ICL port channel 64 went down.
- LLDP neighbors for ports 0/30, 0/31, and 0/32 were deleted.
- Interfaces 0/30, 0/31, and 0/32 went down.
- BFD session flapping and flooding (which was blocked)

```
2024/02/10-13:25:00:653070,  
[LOG-1004],267581/69130,,WARNING,b153_L02,LogmessageBFD-1002floodingdetectedandblocked.  
,bfd_timer.c,line:110,comp:bfdd,ltime:2024/02/10-13:25:00:653014
```

- Cluster keep-alive going down and then coming back up

```
2024/02/10-13:25:01:019241,  
[BGP-1006],267582/69131,,INFO,SLX_L02,BGP:NeighborX.X.X.4onVRFVR7DOWN(RcvNotification :  
HoldTimerExpired).,bgp_port.c,line:2028,comp:bgpd,ltime:2024/02/10-13:25:01:0191422024/  
02/10-13:25:01:110402, [MCT-1008],267583/69132,,INFO,SLX8720-32C,ClusterSDI3-FABRIC-  
cluster-1keep-  
aliveisdown.,mct_trace.c,line:467,comp:mct,ltime:2024/02/10-13:25:01:1102842024/02/10-1  
3:25:01:112023, [NSM-4003],267584/69133,DCE,INFO,SLX_L02,ClusterSDI3-FABRIC-  
cluster-1bringupcompleted,nsm_mct.c,line:1970,comp:nsm,ltime:2024/02/10-13:25:01:111939
```

- BGP neighbor downtime and re-establishment

```
2024/02/10-13:25:11:746200, [BGP-1005], 267587/69134,, INFO, SLX_L02, BGP: Neighbor  
X.X.X.4 on VRF VR7 UP (ESTABLISHED). , bgp_port.c, line: 2059, comp:bgpd,  
ltime:2024/02/10-13:25:11:745002
```

- BFD session flap and recovery for the affected neighbor

```
2024/02/10-13:25:13:504031, [BFD-1001], 267588/69135,, INFO, SLX_L02, BFD Session  
UP for neighbor X.X.X.4 on interface ve711., bfd_timer.c, line: 118, comp:bfdd,  
ltime:2024/02/10-13:25:13:503769
```

The RASlog message on SLX_L02 reported a BFD flap between X.X.X.2 and X.X.X.4. The IP address X.X.X.4 corresponds to ve711 on SLX_FOX2 (a compute), which came back up at 13:25:13.

The unexpected downtime of the secondary path from X.X.X.2 to X.X.X.4 contributed to the traffic loss.

XCO or EFA Authentication Stops Working when Users Configure a Different Authentication

The following symptoms are observed when you configure a different authentication method on TPVM:

1. EFA or XCO authentication fails when LDAP (or TACACS+) is configured in TPVM.
2. Upon removing the LDAP (or TACACS+) configuration from TPVM, HOST authentication is used, and EFA or XCO functions properly again.

When LDAP (or TACACS+) is configured in TPVM, users configured in LDAP are identified as HOST users in EFA or XCO. This causes the authentication configured on EFA or XCO to fail.

Resolution

- Add role mapping for the TPVM user, similar to HOST users.

```
$ efa auth rolemapping add --name user1 --role FabricAdmin --type user
```

- After adding this role mapping, the user will appear as both a HOST and an LDAP user.

```
(efa:extreme)extreme@tpvm-76:~$ efa auth rolemapping show
+-----+-----+-----+-----+-----+-----+
+-----+
| ID | Name | Role |
Type | Auth Type | Auth Identifier |
+-----+-----+-----+-----+-----+-----+
| 11 | cn=user1,ou=people,dc=ldap,dc=extreme,dc=in | FabricAdmin
| USER | LDAP | ldap1 |
+-----+-----+-----+-----+-----+-----+
| 9 | user1 | FabricAdmin
| USER | HOST |
+-----+-----+-----+-----+-----+-----+
| 1 | extreme | SystemAdmin
| USER | HOST |
+-----+-----+-----+-----+-----+-----+
```

- Now, login should work with LDAP configuration present in TPVM.

```
(efa:extreme)extreme@tpvm-76:~$ efa login --username user1 --password userlabc

Login successful.
--- Time Elapsed: 170.602251ms ---
(efa:user1)extreme@tpvm-76:~$ efa inventory device execute-cli --ip 10.20.62.206 --command "show
run tpvm | include ldap"
Execute CLI[success]
+-----+-----+-----+-----+-----+-----+
+-----+
| IP Address | Host Name | Fabric | Command | Status | Reason |
|
| Output |
+-----+-----+-----+-----+-----+-----+
| 10.20.62.206 | SLX | | show run tpvm | include ldap | Success | | SLX# show
run tpvm | include
ldap
|
| | | | | | | | ldap host
10.20.61.63 port 389 basedn dc=ldap,dc=extreme,dc=in rootdn cn=admin,dc=ldap,dc=extreme,dc=in
rootdnpw $9$BwrsDbB+tABWGWpINOVKoQ== |
| | | | | | |
|
+-----+-----+-----+-----+-----+-----+
+-----+
Execute CLI Details
--- Time Elapsed: 5.375033382s ---
```

EFA CLI Login Fails with Error "Failed to Fire Hook"

The issue affects both local and LDAP users when you:

1. Create a new local user on TPVM and assign an XCO role.
2. Log in to TPVM as the new user and attempt EFA CLI login.

```
testuser@b158-12-tpvm:~$ efa login

ERRO[0000] Failed to open log file: open /apps/efa_logs/efa-client.log: permission
denied

Failed to fire hook: invalid argument
Failed to fire hook: invalid argument
Failed to fire hook: invalid argument

Password:
Failed to fire hook: invalid argument
Login successful.

--- Time Elapsed: 3.2229171s ---
```

The error occurs due to changed permissions on `/apps/efa_logs/efa-client.log` (from 666 to 660), requiring `efa` group membership. However, adding users to the `efa` group raises security concerns, as they can access sensitive files like `/apps/efadata/misc/database/dbconfig.properties`, which contains the root DB password.



Note

When `efa-client.log` file is missing, the system attempts to create it during XCO operations and set required privileges. If the current logged-in user lacks necessary permissions, the operation fails. Users must ensure they have adequate privileges.

If the file exists with incorrect ownership (not `root:efa`), delete it and retry with elevated privileges to create it correctly.

Resolution


Complete the following workaround:

- Manually restore the correct permissions:

```
sudo chown efa:efa /apps/efa_logs/efa-client.log
sudo chmod 664 /apps/efa_logs/efa-client.log
```

- Or reset ownership recursively for the full log directory:

```
sudo chown -R efa:efa /apps/efa_logs
```



Troubleshooting Installation and Deployment

- [Connection to Server Localhost: 8080 Fails](#) on page 77
- [EFA or XCO OVA Image Reports Read Only File System](#) on page 78
- [EFA or XCO Deployment Fails on TPVM](#) on page 78
- [EFA or XCO Does Not Start Properly and the PODs Stuck in Init State](#) on page 79
- [Identify the Active Node that Serves as the Database for Kubernetes Clusters](#) on page 80
- [Installation of Database \(MariaDB\) Fails](#) on page 80
- [Installation of XCO or EFA Fails](#) on page 80
- [Login to XCO CLI Fails](#) on page 80
- [Pod is in a Crashloopback State](#) on page 81
- [Post-Uninstallation of EFA or XCO, Existing Fabric Configuration Persists on SLX](#) on page 81
- [RabbitMQ Pods Continuously being Deleted and Redeployed or in CrashLoopBackOff](#) on page 82
- [Split Brain Issue in EFA or XCO HA Deployment](#) on page 83
- [TPVM Deployment Fails](#) on page 84
- [Transport Endpoint is Not Connected](#) on page 85
- [Uninstall a Failed or Partial Installation](#) on page 86
- [XCO Deployment with Management Sub-Interface Fails](#) on page 86
- [XCO Instability After Configuring TPVM LDAP Host](#) on page 87

Learn about troubleshooting installation and deployment issues.

Connection to Server Localhost: 8080 Fails

When a connection to the server at localhost: 8080 fails, complete the following steps:

1. Verify the K3s service status.
 - Ensure that the K3s service is up and running. See [Logging and Log Files](#).
2. Review K3s logs.
 - Check the journalctl log for specific error messages
3. Verify the gateway reachability in a multi node deployment.

Ensure that the gateway is reachable from the XCO host.

EFA or XCO OVA Image Reports Read Only File System

You have installed EFA or XCO on a virtual machine (VM) with storage mounted on a separate ESXi host. A network interruption caused a loss of access to the storage host, making EFA or XCO inaccessible.

The issue occurred because the EFA or XCO file system has parts mounted with the `errors=remount-ro` option, which remounts them as read-only when write errors occur.

Resolution

1. Reestablish connectivity to the storage location.
2. Perform a File System Consistency Check (FSCK) to identify and fix any file system issues.
 - The VM will reboot during this process.

Once the VM restarts, EFA or XCO must be accessible again.

EFA or XCO Deployment Fails on TPVM

The following CLI output indicates that an EFA or XCO installation on TPVM has failed:

```
# efa deploy --graphics no
Starting "efa deploy"...
Step 1: Checking if TPVM is deployed...
TPVM is not installed or configured with right options.

Please clear any warnings/errors with TPVM startup.
The output of 'show tpvm status' should be as follows:
SSH and Sudo passwordless :Enabled
AutoStart :Enabled
Tpvm status :Running
Any warnings can be cleared using show tpvm status clear-tag <tag>
After clearing all error/warning messages run the 'efa ' command again
Use the TPVM Command to deploy tpvm
tpvm deploy [IF_OPTION] [NOPASS_OPTION] [ADMIN_PWD_OPTION]
where:
IF_OPTION:= interface { mgmt | insight } ip-addr { dhcp | static-ip }
static-ip := <ip-addr/netmask> [gateway]
NOPASS_OPTION := passwordless
ADMIN_PWD_OPTION := <enter new password> <confirm new password>

Please choose: 1 Single-node deployment 2 Multi-node deployment
2
Multi-node Deployment
Enter a list of peer nodes (IP/Hostname) separated by spaces:
x.x.x.x
Verifying connectivity to x.x.x.x...
x.x.x.x server is not reachable...
Please enable password-less SSH access between the nodes
The installation cannot proceed
Non-interactive ssh connection to x.x.x.x did not succeed. Please try again.
```

The EFA or XCO deployment logs indicate a permissions issue.

```
cat /var/log/efa_deploy.log
INFO:root:['Aborted: permission denied']
INFO:root:aborted:permissiondenied
```

Resolution

To successfully deploy EFA or XCO on TPVM, complete the following steps:

1. Ensure TPVM is properly configured.
 - Check TPVM status.


```
show tpvm status
```
 - Clear any warnings or errors.


```
show tpvm status clear-tag <tag>
```
 - Verify SSH and Sudo passwordless access is enabled, AutoStart is enabled, and TPVM status is running
2. Deploy TPVM using the TPVM command.


```
tpvm deploy [IF_OPTION] [NOPASS_OPTION] [ADMIN_PWD_OPTION]
```

 - When prompted, choose deployment type: single-node (1) or multi-node (2)
 - For multi-node deployment, enter peer nodes (IP/Hostname) and ensure password-less SSH access between nodes.
3. Resolve permissions issues.
 - Switch from TACACS user to local user.
 - Update AAA authentication to use local-auth-fallback as a backup authentication.

EFA or XCO Does Not Start Properly and the PODs Stuck in Init State

Running the **efactl status** command from the XCO command line shows that the EFA services are stuck in the Init status.

```
(efa:extreme)extreme@tpvm22:~$ efactl status
Node: efa
NAME                READY  STATUS   RESTARTS   AGE    IP        NODE
NOMINATED NODE     READINESS GATES
pod/gofabric-service-s6t4g    0/1    Init:1/2    0         6d16h   1.1.1.1   efa
<none>                <none>
[truncated]
```

This issue occurs due to the installation of supplemental packages on the Ubuntu server.

Resolution

Ensure that any Ubuntu server where EFA or XCO is installed does not have any other (supplemental) packages installed.



Note

For successful installation of EFA or XCO, try removing the Docker service using the `snap remove docker`.

Identify the Active Node that Serves as the Database for Kubernetes Clusters

To identify the active node that serves as the database for Kubernetes clusters.

1. Run the `ip addr show` command on all nodes.
2. Verify that on one of the Ethernet interfaces, the virtual IP address shows up as the secondary IP address.

Installation of Database (MariaDB) Fails

Review the installation log files to identify the reason of failure. The installer logs specify the errors that occurred during the installation process. Additionally, check the deb installation packages in the same log for any issues.

You can also use the `journalctl -u mariadb` command to review the MariaDB journal log.

During installation or upgrade of XCO, the `efamonitor` system service is set up to run validation every minute to ensure the XCO database cluster, glusterFS, and RabbitMQ are functioning correctly.

The `efamonitor` service can remediate issues with the MariaDB Galera cluster and RabbitMQ connection, and logs the system statistics.

Installation of XCO or EFA Fails

Review the installation log files to identify the reason of failure. The installer logs specify the errors that occurred during the installation process.

Use the `efa version` command to see the installation details. If the EFA binary is not available, check the config file:

- For TPVM: `/apps/etc/efa/efa.config`
- For Server: `/etc/efa/efa.config`

For more information on XCO log locations, see [Logging and Log Files](#) and [XCO Installer Improvements for Server-Based Deployment](#).

Login to XCO CLI Fails

If you encounter login failures to XCO CLI, refer to the following solutions:

- **Services Not Operational Error**

```
$ efa login
Password:
CLI Failed.
EFA services are not operational yet.
```

```
Please check service state using 'efactl status' command and retry the login once the
services are up.
```

If you receive a "services not operational" error, complete the following steps:

- Verify the status of XCO deployment by running the XCO **efactl status** script (or the **efa status** command, as an alternative).
- Verify that all PODs are in a running state by running the **k3s kubectl get pods -n efa** command.

- **Invalid Client ID Error**

```
$ efa login
Password:
CLI Failed.
Error Message: client ID is not valid
```

If you see an "invalid client ID" error, update the environment by running the **source /etc/profile** command or by opening a new shell. If issues persist, restart the goauth service.

- **Invalid Credentials Error**

```
$ efa login
Password:
CLI Failed.
Error Message: invalid credentials.
```

If you see an "invalid credentials" errors, complete the following steps:

- Verify that the entered user credentials are correct.
- Ensure that the authentication service is running by using the **systemctl status hostauth.service** command.

Pod is in a Crashloopback State

To know if a POD is in a crashloopback state, run the **k3s kubectl get pods -n efa** command.

To view the logs, run the **k3s kubectl -n efa logs pod/<pod-name-here>** command.

Also see [Avoiding CrashLoopBackOff State for CoreDNS](#)

Post-Uninstallation of EFA or XCO, Existing Fabric Configuration Persists on SLX

When you uninstall EFA or XCO, the fabric configuration remains on SLX devices. This is an expected behavior.

Uninstalling only stops services and removes databases, but doesn't delete tenant, fabric, and inventory devices.

Resolution

To completely uninstall EFA or XCO and remove fabric configuration from SLX devices, complete the following steps:

1. Delete all tenant services.

```
efa tenant delete --name [tenant_name] --force
```

The `--force` option automatically removes any underlying tenant objects such as EPG, VRF, and PO.

2. Delete fabric from EFA or XCO.

```
efa fabric delete --name [fabric_name] --force
```

The `--force` option automatically removes all devices from the fabric along with the EFA managed fabric configuration.

3. Delete devices from EFA or XCO inventory.

```
efa inventory device delete --ip [ip_address|range]
```

This will remove any inventory settings such as SNMP server configurations.

4. Uninstall EFA or XCO.

- For TPVM based deployments, run the following command from the SLX:

```
no efa deploy
```

For more information, see [Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#).

- For Server based deployments, run the following command:

```
source ./deployment.sh -o undeploy
```

For more information, see [Uninstall XCO in a Single-Node or Multi-Node Deployment](#).

RabbitMQ Pods Continuously being Deleted and Redeployed or in CrashLoopBackOff

The persistent volumes for RabbitMQ pods reduces message loss during double faults or system failures. These persistent volumes are tightly linked to the k3s host and the specific RabbitMQ pod deployed on that host. Pre-existing files on these volumes containing configuration data for an existing RabbitMQ cluster can affect the formation of new clusters.

During installation, upgrade, restore, or double faults (or system failures), RabbitMQ pods might be assigned to different k3s hosts. If old configuration data is present on these persistent volumes, and if it is not updated, it can lead to issues with new cluster formation. If this occurs, the RabbitMQ pods fail to form a cluster.

The monitoring service in XCO detects and remediates this condition. The remediation script tries several simple fixes to minimize message loss before resetting the cluster. These fixes may result in multiple deletions and redeployments of RabbitMQ pods.

You can monitor the progress of these remediation processes by reviewing the logs located at the following locations:

- For TPVM: `/apps/efa_logs/efa-monitor/rabbitmq*.log`
- For Server: `/var/log/efa/efa-monitor/rabbitmq*.log`

Generally, XCO services remain functional during remediation as the remediation script attempts to sync with the RabbitMQ pod on the standby node. In some cases, XCO might take longer to come online. No additional actions are required from users.

Split Brain Issue in EFA or XCO HA Deployment

When upgrading EFA or XCO in an HA setup, the following issues are observed:

- Slow response from EFA
- VIP bouncing between both TPVMs
- Attempts to change the TPVM hostname timeout after 30 minutes

This is often caused by using capital letters in the TPVM hostname, which Kubernetes doesn't support.

Resolution

Complete the following steps:

1. Record the old and new hostname values and identify the standby TPVM.

```
EFA# efa status
```

2. Stop the EFA monitor service on the standby TPVM via SSH.

```
TPVM# sudo systemctl stop efamonitor
TPVM# exit
```

3. Change the TPVM hostname to lowercase via SLX CLI via SSH.

```
SLX-SWITCHX# tpvm config hostname <new hostname following the recommendation>
```

4. Run the hostname change script on the standby TPVM.

```
TPVM# sudo /apps/bin/efa-change-hostname <OLD TPVM HOST NAME>
```

The standby TPVM will go down during this process. The script will report `restartmariadb, k3s, and so on, and finally, Successfully updated hostname in EFA.`

Wait for EFA HA stability (the `efa status` must show both active and standby as “up” with the new hostname visible on the updated TPVM).

5. Restart the EFA monitor service and log out of TPVM.

```
TPVM# sudo systemctl start efamonitor
TPVM# exit
```

6. Validate the hostname in the `mgmt_networks.txt` file.

```
TPVM# cat /apps/efadata/misc/multiaccess/mgmt_networks.txt
```

If the hostname is still in uppercase, correct it manually.

7. Ensure keepalived configurations match on both devices.

```
TPVM# cat /etc/keepalived/keepalived.conf
```

Ensure that both files show the same VIP and that the number of configured networks matches.

8. Trigger an EFA or XCO HA failover and repeat steps 2-6 for the remaining TPVM.

```
SLX-SWITCHX# tpvm stop
SLX-SWITCHX# tpvm start
```

9. Wait for EFA HA stability (both active and standby must be up), then repeat steps 2 to 6 for the (previously active, now standby) TPVM.

TPVM Deployment Fails

The TPVM deployment fails due to a configuration issue.

TPVM Deployment Error: TPVM Trusted Peer Configuration Failed

The following commands were run to uninstall the TPVM:

```
INFO Executed command: 'tpvm stop'
INFO Executed command: 'tpvm uninstall force'
INFO Executed command: 'configure terminal'
INFO Executed command: 'tpvm TPVM'
INFO Executed command: 'no deploy'
INFO Executed command: 'firmware download fullinstall scp directory /var/lib/hds/repo/slx/slx-os/current/slxos20.4.1ca host 21.151.151.254 user operator password'
INFO Executed command: 'copy default-config startup-config'
INFO Executed command: 'reload system'
INFO Executed command: 'configure terminal'
INFO Executed command: 'tpvm TPVM'
INFO Executed command: 'no deploy'
INFO Executed command: 'password xxxxxx,30'
INFO Executed command: 'auto-boot'
INFO Executed command: 'ntp 21.151.151.254'
INFO Executed command: 'hostname b151-101'
INFO Executed command: 'timezone Europe/Stockholm'
INFO Executed command: 'interface management ip 21.151.151.2/24 gw 21.151.151.254'
INFO Executed command: 'deploy'
INFO Command Response:
'% Error: TPVM trusted peer configuration failed.'
```

The issue arises from a script that was used to uninstall and reinstall TPVM. Specifically:

```
INFO Executed command: 'no tpvm TPVM'
```

- This command removes the TPVM configuration from the running-config, but the configuration remains in the startup-config.
- After this, the script performs the **firmware download fullinstall** and reloads the system. Upon reload, the configuration is replayed from the startup-config, reintroducing the TPVM trusted-peer configuration and other TPVM config back into the running-config.
- In this state, when the script runs the **copy default-config startup-config** command, the TPVM configuration is persisted.
- When the script tries to deploy TPVM again, it fails due to the persisted configuration.

Resolution

Add `running-config startup-config` to the script immediately after `no tpvm TPVM` or use `copy default-config startup-config remove-tpvm` after the fullinstall.

Transport Endpoint is Not Connected

The `efamonitor` service on the standby node has been down since a long time (for example, January 24th).

```
efamonitor.service - EFA node and service monitor service
Loaded: loaded (/lib/systemd/system/efamonitor.service; enabled; vendor preset: enabled)
Active: inactive (dead) since Wed 2024-01-24 07:21:52 GMT; 2 weeks 2 days ago
Process: 694 ExecStart=/apps/bin/efa-monitor.sh (code=killed, signal=TERM)
Main PID: 694 (code=killed, signal=TERM)
```

Any attempt to restart the `efamonitor` service result in a continuous error message: "Transport endpoint is not connected":

```
Feb 16 17:07:45 xxxxxxxx efa-monitor.sh[11112]: stat: cannot stat '/apps/efa_logs':
Transport endpoint is not connected
```

The issue occurred because the `glusterfs` mount failed, leading to the "Transport endpoint is not connected" error and causing the `efamonitor.service` to become inactive.

Resolution A

Upgrade to XCO 3.2.0 or above.

Alternatively, reload the standby TPVM to remount `glusterfs` correctly.

Resolution B

To reestablish the transport endpoint connectivity, complete the following steps:

1. Verify the status of `glusterfs` on both nodes in the high availability (HA) setup.
2. Check if the following three mount points are available:
 - a. `/apps/efadata/certs`
 - b. `/apps/efa_logs/`
 - c. `/apps/efadata/misc`
3. If the mount points are present, perform the following steps on the node where the "Transport endpoint is not connected" error is occurring:
 - a. Restart the `glusterd` service.

```
sudo systemctl restart glusterd.service
```
 - b. Restart the monitoring service.

```
sudo systemctl restart monitor.service
```

Uninstall a Failed or Partial Installation

For a failed or partial installation of XCO, run the following command to uninstall or unwind the XCO instance:

- For a TPVM based deployments - **no efa deploy**

For more details, see [Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#).

- For a Server based deployments - **source deployment.sh**

For more details, see [Uninstall XCO in a Single-Node or Multi-Node Deployment](#).

XCO Deployment with Management Sub-Interface Fails

During XCO deployment, the following command fails when trying to add management sub-interface:

```
efa deploy non-interactive multi-node package /efaboot/efa-3.3.1.tar.gz peer-node
10.152.3.40 vip4 10.152.3.41 management-ip sub-interface-name ext-EFA sub-vlan-id 4054
external-subnet 10.33.57.197/28
```

The deployment of XCO fails because of a duplicate IP address or route configuration on the Eth0 network interface.

Resolution

1. Check the EFA logs (installer debug.log).

```
installer_192.168.246.3_20231030_debug.log:
=====
Target_External_Networks=network 10.33.45.197/28 via interface ext-EFA on vlan 4054 <
Target_Release=3.3.1-GA
Target_Version=3.3.1

+(2023-10-30T14:21:36.473010 GMT b147-11 common.sh:1797): change_default_route_interface():
rt='default via 10.33.57.195 dev eth0 '
+(2023-10-30T14:21:36.485844 GMT b147-11 common.sh:1800): change_default_route_interface(): [[ -z
default via 10.33.57.195 dev eth0 ]]
+(2023-10-30T14:21:36.502170 GMT b147-11 common.sh:1806): change_default_route_interface(): ip -4
route add default via 10.33.57.195 dev eth0 table 99
+(2023-10-30T14:21:36.528757 GMT b147-11 common.sh:1808): change_default_route_interface(): false
+(2023-10-30T14:21:36.550226 GMT b147-11 common.sh:1810): change_default_route_interface(): ip -4
route delete default via 10.33.57.195 dev eth0
+(2023-10-30T14:21:36.568956 GMT b147-11 common.sh:1815): change_default_route_interface(): false
++(2023-10-30T14:21:36.582826 GMT b147-11 common.sh:1818): change_default_route_interface():
change_route_interface 'default via 10.33.57.195 dev eth0 ' efa-ext-EFA
++(2023-10-30T14:21:36.593075 GMT b147-11 common.sh:1782): change_route_interface(): local
'rt=default via 10.33.57.195 dev eth0 '
++(2023-10-30T14:21:36.602698 GMT b147-11 common.sh:1783): change_route_interface(): local
if2=efa-ext-EFA
++(2023-10-30T14:21:36.618817 GMT b147-11 common.sh:1784): change_route_interface(): printf
'%s\n' 'default via 10.33.57.195 dev eth0 '
++(2023-10-30T14:21:36.621205 GMT b147-11 common.sh:1784): change_route_interface(): sed -E
's/dev [^ ]+ /dev efa-ext-EFA /'
++(2023-10-30T14:21:36.621652 GMT b147-11 common.sh:1784): change_route_interface(): sed -E
's/dev [^ ]+$/dev efa-ext-EFA/'
```

The EFA or XCO attempts to add and delete the route on eth0 but fails, then tries to change the route on the sub-interface and fails again.

2. Check the keepalived logs (/keepalived/journalctl_IP.log).

```
Oct 30 16:53:04 b147-12 Keepalived_vrrp[911197]: (/etc/keepalived/keepalived.conf: Line 39)
WARNING - interface efa-ext-EFA for ip address 10.33.57.197/28 doesn't exist
```

In this scenario, the default route overlaps with the sub-interface network. The system keeps trying to change the default route, but since the sub-interface does not have an IP address, it continuously fails.

3. Verify and confirm the route overlap in the installer logs.

```
{"@time":"2023-10-30T14:21:34.028321 GMT","App":"common","msg":"Instantiating management interface
efa-ext-EFA"}
{"@time":"2023-10-30T14:00:13.571263 GMT","App":"common","msg":"sub_intfname: ext-EFA "}
{"@time":"2023-10-30T14:00:13.599507 GMT","App":"common","msg":"sub_vlanid : 4054"}
{"@time":"2023-10-30T14:00:13.627877 GMT","App":"common","msg":"cidr : 10.33.57.197/28 "}
```

The same IP existed before installation, confirming the route overlap issue.

XCO Instability After Configuring TPVM LDAP Host

After configuring LDAP on TPVM, the standby XCO node experiences instability, repeatedly going up and down. However, the XCO on the active node functions properly with successful login, command execution, and all the K3s PODs running smoothly.

The standby node faces issue due to the following reasons:

- SSH connection failures between active and standby nodes, preventing status collection.

```
/apps/efa_logs/efact1/efact1_10.53.2.21_debug.log

+++ (2023-12-26T12:57:10.320365 PST tpvm21 common.sh:23697): efa_node_status_json():
efa_node_status_to_json tpvm22 down 10.53.2.22 standby 0 false 'Cannot connect with
ssh' 0 0 0
```

- LDAP connection errors on the standby node, including failed binds and server unavailability.

```
journalctl -u ssh

Dec 26 12:57:30 tpvm22 sshd[71891]: nss_ldap: could not connect to any LDAP server as
(null) - Can't contact LDAP server
Dec 26 12:57:30 tpvm22 sshd[71891]: nss_ldap: failed to bind to LDAP server ldap://
10.10.10.180:389/: Can't contact LDAP server
Dec 26 12:57:30 tpvm22 sshd[71891]: nss_ldap: reconnecting to LDAP server...
```



Note

This issue occurs when management sub-interfaces and routes are configured only on the active node, making the LDAP server inaccessible to the standby node.

Resolution


Upgrade to TPVM 4.6.7 or above.

Alternatively, complete the following workaround:

1. Update the `/etc/ldap.conf` file on TPVM with the following specified settings:

```
nss_initgroups_ignoreusers _apt,backup,bin,daemon,efainternal,extreme,fwupd-  
refresh,games,gluster,gnats,irc,landscape,list,lp,lxd,mail,man,messagebus,mysql,news,no  
body,pollinate,proxy,root,_rpc,sshd,statd,syslog,systemd-coredump,systemd-  
network,systemd-resolve,systemd-timesync,tcpdump,tss,usbmux,uucp,uidd,www-data  
bind_policy soft
```

2. Ensure both active and standby TPVM nodes use a locally accessible LDAP server.



Troubleshooting Network Infrastructure Components

[RabbitMQ Log Rotation Does Not Remove or Archive Old Logs](#) on page 92

[Two Instances of Each Service is Seen in EFA or XCO](#) on page 92

[Verify Reachability of XCO to a New Management IP Address](#) on page 93

Learn about troubleshooting issues with network infrastructure components.

RabbitMQ Log Rotation Does Not Remove or Archive Old Logs

In EFA 3.0.0, RabbitMQ logs do not rotate properly, leading to excessive disk space usage.

To overcome the issue, upgrade to EFA 3.0.1, XCO 3.2.0 or above. Alternatively, complete the following steps:

Resolution

Upgrade to EFA 3.0.1, XCO 3.2.0 or above.

Two Instances of Each Service is Seen in EFA or XCO

High disk utilization (above 93%) in EFA or XCO triggers a restart, resulting in duplicate service instances. However, the restart fails to remove old instances, causing conflicts due to outdated database content.

This issue affects services like auth, fabric, and tenant, with old tenant instances retaining outdated database content. Error logs indicate image garbage collection failed due to high disk usage (93%), and stats initialization may be incomplete.

```
Jul 26 12:32:08 seroius12382 k3s[1129]: I0726 12:32:08.177745 1129
image_gc_manager.go:305] [imageGCManager]: Disk usage on image filesystem is at 93% which
is over the high threshold (85%). Trying to free 6401413939 bytes down to the low
threshold (80%).
Jul 26 12:32:08 seroius12382 k3s[1129]: E0726 12:32:08.180089 1129 kubelet.go:1321]
Image garbage collection failed once. Stats initialization may not have completed yet:
failed to garbage collect required amount of images. Wanted to free 6401413939 bytes, but
freed 0 bytes
efa_gofabric-service-cbndn_af269635-f218-48ff-8e69-b92ec85a035d/ 2023-Jul-27
09:11:13 - Directory
efa_gofabric-service-gdrbz_e7887d3a-c756-4e07-8431-99ab41cfdb4e/ 2023-Jul-27
```

```

09:11:13 - Directory
fabric_database_dump_2023-06-22T12-51-47.524.sql 2023-Jul-27
09:11:13 63.5M application/x-sql
fabric_database_dump_2023-07-27T12-23-46.773.sql 2023-Jul-27
09:11:13 64.7M application/x-sql

efa_gotenant-service-4848r_71b4e6f3-b308-4b5c-8e31-f097bcca61ce/ 2023-Jul-27
09:11:30 - Directory
efa_gotenant-service-tdqt7_3f6613f4-f315-402b-8524-28eef98ee0cd/ 2023-Jul-27
09:11:29 - Directory
ts_database_dump_2023-06-22T12-53-58.116.sql 2023-Jul-27
09:11:30 37.6M application/x-sql
ts_database_dump_2023-07-27T12-23-48.854.sql 2023-Jul-27
09:11:29 40.8M application/x-sql
    
```


Resolution

To recover from the high disk usage, complete the following steps:

1. Clear disk space and restart the services to ensure only new instances run.
2. Run the **efactl clean** to clean up disk space.
3. Re-import images using the `k3s ctr image import /opt/efa/docker_images/docker_k3s_images.tar` file.
4. Restart EFA or XCO and k3s services using the **efactl restart** and **systemctl restart k3s** commands to restore normal functionality.

Verify Reachability of XCO to a New Management IP Address

To confirm that the XCO server is reachable from a new management IP address, make sure the VLAN configuration of the sub-interface and the IP address subnet are compatible with the host attempting to establish a connection (ping) to the XCO server.



Troubleshooting Policy Services

[QoS Profiles and Maps are not Automatically Removed](#) on page 94

Learn about troubleshooting policy service issues.


QoS Profiles and Maps are not Automatically Removed

QoS profiles remain bound, and QoS maps are not removed when a fabric or tenant is deleted.

This is an expected behavior, as the system does not automatically unbind QoS policies during deletion. Before performing any force operations, such as deletions, ensure that QoS policies are unbound from all associated targets, including fabrics, tenants, ports, port channels, and tenant endpoint groups, to prevent stale policies from lingering in the system.

Resolution

Manually unbind the QoS policies before deleting the fabric, tenant, port, port channel, and tenant endpoint group.



Troubleshooting Tenant Services

[Adding Multiple BFD and Static Routes to SLX Config via XCO Causes a Panic on SLX on page 95](#)

[APS Admin Down or Admin Up Operations Fail on page 96](#)

[BGP Peer Group Creation Fails Due to Permanent EFA Error on page 96](#)

[Change Anycast MAC Address in SLX Fabric Without Impacting Services on page 97](#)

[EPG Creation or Update Fails on page 97](#)

[EFA or XCO Tenant Command Fails on page 98](#)

[Failed to Save Devices During EPG Creation on page 98](#)

[How to Enable BFD Under a VE Interface via XCO on page 99](#)

[MCT Cluster Config Missing Between Rack Devices Error Message is Not Relayed in REST API on page 100](#)

[Static Routes Persist on SLX Devices after Deletion from XCO on page 101](#)

[Successful Tenant Creation with Overlapping L2 and L3 VNIs on page 101](#)

[Tenant Creation Fails on page 102](#)

[Tenant Service Restart Scenario on page 103](#)

[Tenant VRF Update Fails on page 105](#)

[Unable to Delete an EPG with Comma in Name on page 106](#)

[VLAN Configuration Fails When Adding Port or PO to EPG on page 107](#)

[VRF Already Exists on Device Error on page 108](#)

Learn about troubleshooting tenant service issues.

Adding Multiple BFD and Static Routes to SLX Config via XCO Causes a Panic on SLX

When adding static routes via XCO, the associated border-leaf devices unexpectedly reloads due to the SRMD process.

Adding and deleting BFD routes causes a memory leak in the static route space. The memory leak eventually leads to the SRM (Static Route Manager) daemon encountering the leaked memory space, resulting in an unexpected reload.

Resolution

Upgrade to SLX-OS 20.6.1 or later.

Alternatively, after adjusting BFD routes, reload the device to clear the memory space.

APS Admin Down or Admin Up Operations Fail

Use this topic to learn how to resolve issues related to admin down or admin up operation failures.

APS Admin Down Operation Fails

1. Run the **efa inventory admin-state detail --uuid** command to check the detailed status of the fabric and tenant.
2. If any service status is reported as failed, use the UUID to check the corresponding service logs and identify the reason for the failure.

APS Admin Up Operation Fails

1. Run the **efa inventory admin-state detail --uuid** command to check the DRC status.
2. If the DRC status is failed, run the **efa inventory drift-reconcile detail --uuid** command to find the specific cause of the DRC failure.
3. If the DRC status is successful, review the status of the fabric or tenant services.
4. If any of the fabric or tenant services have failed, use the UUID to verify the respective service logs and determine the cause of the failure.

BGP Peer Group Creation Fails Due to Permanent EFA Error

Attempting to create a BGP peer group resulted in the following error:

```
efa tenant service bgp peer-group create --name BL-VR1-DBGP-PG -- tenant Tnt-1 --pg-name 21.130.130.201:VR 0.201, VR1-PeerGrp: 1000 - -pg-bfd-enable 21.130.130.201, VR1-PeerGrp: true - -pg-bfd 21.130.130.201, VR1-PeerGrp: 130.202:VR1-PeerGrp --pg-asn 21.130.130.202, VR1-PeerGrp: 1000 --pg-bfd-enable 21.130.130.202, VR1-PeerGrp: tR1-PeerGrp: 300,300,3
```

```
Error : Error 1452 (23000): Cannot add or update a child row: a foreign key constraint fails ('dcapp_ten ibutes', CONSTRAINT "bp_peer_group_afi_attributes_ibfk_2" FOREIGN KEY ('af_id') REFERENCES 'bgp_address_ADE)
```

Resolution

Delete the tenant using the `--force` option.

Change Anycast MAC Address in SLX Fabric Without Impacting Services

In a network setup, two sites shared the same anycast MAC ("0201.0102.0102") for a specific VLAN XXX, causing traffic loss.

Follow this procedure to update the anycast MAC address in an SLX fabric without disrupting services.



Note

Ensure that you run this procedure under the supervision of SME, ESCL, or GTAC engineers

1. Update the anycast MAC in the `fabric_properties` table of the fabric database.

```
UPDATE fabric_properties
SET any_cast_mac="new_mac_address",
ip_v6_any_cast_mac="new_ipv6_mac_address"
WHERE id=fabric_id;
```

2. Remove existing anycast IP addresses from devices using the following command:

```
efa inventory device execute-cli
--ip device_ip
--command "no ip anycast-gateway-mac old_mac"
--config

efa inventory device execute-cli
--ip device_ip
--command "no ipv6 anycast-gateway-mac old_ipv6_mac"
--config
```

3. Perform drift and reconciliation (DRC) to apply changes.

```
efa inventory drift-reconcile execute
--ip device_ip
--reconcile
```

4. Confirm the updated anycast MAC addresses using the following command:

```
efa inventory device execute-cli
--ip device_ip
--command "show running-config | inc anycast"
```

EPG Creation or Update Fails

When creating or updating an endpoint group (EPG), you encounter an error with the code `{'NETWORK': 80}` and the message `Device(s) are locked with reason configbackup`, it's likely due to a scheduling conflict.

```
Error Code: {'NETWORK': 80, 'message': 'Response code received b'
{'message': "Device(s) 10.157.186.28 are locked with reason configbackup", "code":
1002}'}
```

The issue occurred because the default system backup for EFA or XCO systems occurs daily at 00:00:00 hours, and API operations cannot run simultaneously.

Resolution

Ensure that your API operations do not overlap with the default EFA or XCO system backup operation schedule.

EFA or XCO Tenant Command Fails

The **efa tenant** command fails with the following error:

```
http://localhost/v1/tenant/tenants: dial tcp [::1]:80: connect: connection refused
```

The x.509 certificate had expired due to an incorrect date and time settings on the server.

Resolution

Correct the server date and time settings or configure Network Time Protocol (NTP).

Failed to Save Devices During EPG Creation

Attempting to create an EPG results in the following error:

```
Error: EPG: [EPG_Name] Save for Evpn BD Records
```

The error occurs because a stale database entry is conflicting with the desired configuration. According to the "evpn_bridge_domain_mapping" table, the combination of keys (evpn_id, bridge_domain_value) must be unique. An attempt to add another entry with the same key (evpn_id, bridge_domain_value) results in the following error:

- evpn_id: 293
- bridge_domain_value: 118

The following is an example output of the **ts-server_err.log**:

```
{"@time":"2022-02-17T16:33:43.099571
CET", "App": "TenantService", "level": "error", "msg": "[github.extremenetwork
s.com/Engineering/GoDCApp/GoTenant/src/ts-server/infra/database.
(*DBHandle).FirstOrCreateEvpnBridgeDomainsMap:816] : [rows:0, elapsed:
648.394µs] INSERT INTO evpn_bridge_domain_mapping
(evpn_instance_id,bridge_domain_id,bridge_domain_value,oob_created,dev_s
tate,app_state,reason) VALUES (293,212565,118,false,'not-
provisioned','cfg-ready',''): Error 1062: Duplicate entry '293-118' for
key 'evpn_instance_id'}
```

Resolution

1. Redeploy EFA or XCO.
 - Clean up the database by redeploying EFA or XCO.
2. Update device inventory.
 - If the configuration is present on the SLX device, remove the conflicting configuration, then run the `efa inventory device update` command

Run this command multiple times to ensure the update completes successfully.

- If the configuration is not present on the SLX device, add it, run the `efa inventory device update`, remove the configuration, and repeat the update process.
3. Manually delete or cleanup database.

If the step 1 and step 2 do not resolve the issue, and with Engineering approval, you can manually delete the conflicting database entries using the following SQL commands:

Delete duplicate entries from "overlay_gateway_bridge_domain_vni_mapping" and "evpn_bridge_domain_mapping" tables.

```
deletefrom overlay_gateway_bridge_domain_vni_mapping where overlay_gateway_id =293and
bridge_domain_value =118;
deletefrom evpn_bridge_domain_mapping where evpn_instance_id =293and
bridge_domain_value =118;
```

By removing duplicate entries (rows), the issue will be resolved, and EPG addition will succeed. :

```
(efa:admin)admin@efa-singlenode:~$ efa tenant epg create --tenant "Ten1" --name "EPG1"
--description "Desc1" --type extension --switchport-mode trunk-no-default-native --po
[list_of_PO_Names] --vrf [vrf_Name] --ctag-range 34 --anycast-ip 34:10.10.0.1/28 --
local-ip 34,10.10.8.56:10.10.0.2/28 --local-ip 34,10.10.8.55:10.10.0.3/28
```



Note

As a best practice, escalate stale database issues to GTAC for full diagnosis.

How to Enable BFD Under a VE Interface via XCO

Run the `efa inventory device execute-cli` command.

```
(efa:extreme)extreme@tpvm129:~$ efa inventory device execute-cli --ip="xx.xx.3.127" --config --
command="interface ve 822, bfd interval 300 min-rx 300 multiplier 3"
Execute CLI[success]
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| IP Address | Host Name | Fabric | Command | Status | Reason |
| | | | Output | | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| xx.xx.3.127 | C11b_L01-127 | dcc11b | interface ve 822 | Success | |
| C11b_L01-127(config)# interface ve 822 | | | | |
| | | | bfd interval 300 min-rx 300 multiplier 3 | | |
| C11b_L01-127(config-if-Ve-822)# bfd interval 300 min-rx 300 multiplier 3 | | |
| | | | | | |
| | | | | | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
C11b_L01-127# show run in ve 822
interface Ve 822
 ipv6 address 2001:1b74:480:60e6::2/64
 bfd interval 300 min-rx 300 multiplier 3<====
 no shutdown
 !
C11b_L01-127#
```

The device xx.xx.3.127 is the IP address of the SLX.

MCT Cluster Config Missing Between Rack Devices Error Message is Not Relayed in REST API

While adding multiple devices to the fabric, the error message MCT Cluster Config missing between rack devices is not transmitted through the REST API.

```
(efa:extreme)extreme@tpvm2:~$ efa fabric device add-bulk --name=noncl1 --rack noncl1 --ip
10.67.72.232,10.67.72.233 --username admin --password password
Inventory Device(s) Registration[Success]
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | IP Address | Host Name | Model | Chassis Name | Firmware | Status | Reason |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | 10.67.72.233 | Leaf-2 | 4000 | BR-SLX9540 | 20.3.4ab | Device with IP: 10.67.72.233 already registered in the | |
| | | | | | | application | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 3 | 10.67.72.232 | Leaf-1 | 4000 | BR-SLX9540 | 20.3.4ab | Device with IP: 10.67.72.232 already registered in the | |
| | | | | | | application | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Device Details
Updating devices that are already registered: [10.67.72.233 10.67.72.232]
Inventory Update with ip-address = 10.67.72.232 [Succeeded]
Inventory Update with ip-address = 10.67.72.233 [Succeeded]
Add Device(s) [Success]
Addition of Leaf device with ip-address = 10.67.72.233 [Succeeded]
Addition of Leaf device with ip-address = 10.67.72.232 [Succeeded]
Validate Fabric [Failed]
Config MisMatch
MCT Cluster Config missing
between rack devices 10.67.72.233 and 10.67.72.232.Please remove rack devices and re-add.
Error : fabric validation failed
```

However, when attempting to add bulk devices using the EFA CLI, the error message appears, but it's not included in the REST response.

```
gtac@exodus:~$ curl -kX POST 'https://10.67.72.239/v1/fabric/devices' -i
-d '{"fabric-name":"noncl1","items":[{"ip-address":"10.67.72.232","rack":"noncl1","role":"leaf"}, {"ip-
address":"10.67.72.233","rack":"noncl1","role":"leaf"}]}' -H "Content-type: application/json" -H
"Authorization: Bearer $authtoken"
HTTP/2 200
content-type: application/json; charset=UTF-8
date: Tue, 01 Nov 2022 16:22:32 GMT
x-request-id: 8a253080-flae-4cc7-87ca-8b9ce7dbc90b
content-length: 157

{"items":[{"ip-address":"10.67.72.232","role":"Leaf","rack":"noncl1","error":null}, {"ip-
address":"10.67.72.233","role":"Leaf","rack":"noncl1","error":null}]}gtac@exodus:~$
```

This discrepancy occurs because the EFA CLI and REST API don't have a one-to-one correspondence.

For instance, the EFA CLI's `add-bulk` command performs inventory registration, adds devices to the fabric, and validates the fabric topology. Similarly, certain REST APIs lack equivalent CLI commands, such as `fabric validate`.

Resolution

To ensure robustness, run the `fabric validate` REST API before adding (`fabric device add`) or configuring (`fabric configure`) the fabric using the REST API commands.

Static Routes Persist on SLX Devices after Deletion from XCO

Initially, static routes were configured on the SLX devices via XCO. However, when attempting to reuse the routes after deletion, an error occurred. The error message indicated that deletion of Out of Band (OOB) static routes is not supported.

```
{
  "@time": "2024-04-12T14:39:28.864661 CEST",
  "App": "TenantService",
  "level": "error",
  "msg": "Failed",
  "reason": "The out of band (OOB) created static route are 10.10.10.145-
{2000:1:123:4000::/52, 2011:123:4::5}, Deletion of the OOB Static Routes is not supported"
}
```

The issue occurred because of the communication breakdown between XCO and SLX (or use of the `--force` flag) during route deletion, resulting in routes remaining on SLX. Due to the considerable delay between the initial removal and subsequent inventory updates, the root cause of the issue cannot be precisely determined.

Additionally, since inventory updates occurred after the initial removal attempt, XCO now treats these routes as Out Of Band (OOB) configurations, which cannot be managed via XCO. Creating, editing, or deleting OOB configurations via XCO is not supported.

Resolution

Complete the following workaround:

1. Manually delete the offending configurations either directly on the SLX or using the `eexecute-cli` command.
2. Refresh the SLX configuration within XCO for comparison using the `efa inventory device update` command.
3. Retry the configuration change.

Successful Tenant Creation with Overlapping L2 and L3 VNIs

An EFA or XCO tenant has been created successfully despite having overlapping L2 and L3 VNI parameters.

Multiple tenants without a fabric reference (for example, no assigned ports) can have overlapping VNI values. This is because users can add ports from any of the available fabrics in the future, ensuring a valid configuration. For example, tenants "ten1" and "ten2" can have overlapping VNI ranges if they belong to separate fabrics.

Resolution

Multiple tenants within the same fabric (for example, tenants with ports from a specific fabric) cannot have overlapping VNI values.

Tenant Creation Fails

Use this topic to learn about troubleshooting issues with failed tenant creation scenarios.

Tenant Creation Fails with Port Error

Attempting to create a tenant using ports from a newly added device results in an error, indicating that the device port discovery may still be in progress.

Discovery of device ports may be in progress

Resolution

- Verify the tenant service logs at `/var/log/efa/ts/ts-server.log`.

The logs are rotated in the same directory with time stamps. Each log line contains a set of key-value pair and is in JSON format. All the log lines contain severity level, time stamp, and message text as mandatory keys.

By default, only informational-level log messages are recorded.

- Verify the fabric configuration using the `efa fabric show` command.

This command displays the device state and app state, which should show as 'provisioned' and 'in-sync' respectively for the affected device.

Tenant Creation Fails Intermittently After Initial Fabric Creation: Ports Unavailable

Tenant creation fails sporadically after creating the fabric, with an error indicating that ports are not available.

```
Executed command: 'efa tenant create --name leaf-1-shared-tenant --port
xx.xx.246.11[0/21,0/22] --type shared'
#####23-02-26 01:43:26.942
ASSERT
Assertion Failed: (JcatAssertionError: Command execution not successful)

Error : Ports (xx.xx.246.11-ethernet-0/21,xx.xx.246.11-ethernet-0/22) are not available
in the application. Probable reasons are (i) The fabric is not provisioned on the
device(s) (ii) Discovery of the device ports is in progress
```

Fabric creation is a quick database update, while fabric configuration takes longer due to tasks like interface availability checks and LLDP communications.

Resolution

Add a 3-minute delay before creating a tenant or upgrade to XCO 3.2.0 for improved inventory service. Starting from XCO 3.2.0 onwards, inventory service improvements have sped up this process, particularly step 3 in the following sequence:

1. Fabric service sends an event to the inventory service.
2. Inventory service discovers assets and updates its database.
3. Inventory service sends an event to the tenant service after deep discovery.
4. Tenant service populates its database by sweeping the inventory database.



Note

The delay is due to the microservice design and inherent in the process. Port discovery and tenant database population occur in the background, and you must wait before retrying tenant creation.

Tenant Service Restart Scenario

Use this topic if the Tenant Service restarts repeatedly during execution of the **efa system restore** command.

This indicates:

- Missing readiness checks (for example, RabbitMQ not ready)
- Long-running init tasks (like InventoryDeviceSweep) exceeding probe timeouts
- Underlying dependency failures (DB, Fabric, Inventory, or MessageBus)

Scenario 1: Missing RabbitMQ Readiness Check

Current State

Tenant Service uses the following three init containers:

- **check-db-ready**: Waits for DB readiness.
- **wait-for-inventory-fabric-pod**: Polls `/health` endpoints of Inventory and Fabric services.
- **wait-for-rabbitmq**: No command implemented.



Note

The containers exist in GoTenant and other services, but lack readiness logic.

Issue

- No verification that RabbitMQ is ready before Tenant Service starts.
- This can cause **race condition** during startup, which leads to service crashes or restarts if RabbitMQ is not fully initialized.

Solution

Add readiness check to wait-for-rabbitmq init container

```
- name: wait-for-rabbitmq
  command: ["sh", "-c"]
  args:
    - |
      until wget -qO- http://
      $MessageBusUser:$MessageBusPassword@$MessageBusHost:$MessageBusMgmtPort/api/aliveness-
      test/%2F; do
        echo "Waiting for RabbitMQ to be ready...";
        sleep 5;
      done;
      echo "RabbitMQ is ready!";
```

Before proceeding, Tenant Service now waits until **RabbitMQ is fully operational** to prevent startup failures.

Scenario 2: Liveness Probe Failure During Long Initialization

Issue

- After init containers complete, if Fabric or Inventory services go down, Tenant Service may fail without recovery. Tenant Service depends on Fabric and Inventory being fully up.
- Startup time may exceed probe thresholds due to:
 - Sequential execution of:
 - CleanupInventoryInProgressDevices
 - InventoryDeviceSweep: **Dominant time factor**
 - Sweep duration scales with:
 - Number of managed devices
 - Load on Fabric/Inventory services

Solution

Define a reasonable startup timeout before API Server begins accepting requests.

API Server Startup Sequence

1. DB initialization
2. RabbitMQ initialization
3. Service handler registration
4. CleanupInventoryInProgressDevices
5. Critical Path: InventoryDeviceSweep dominates with startup time and scales with device count and backend load

Configure Startup Probe

To prevent **premature restarts** during long initialization, use a **startupProbe**.

```
startupProbe:
  httpGet:
    path: /v1/tenant/health
    port: 8083
  initialDelaySeconds: 20
```

```
periodSeconds: 20
failureThreshold: 90
```



Note

- periodSeconds: 20: Probe executes every 20 seconds
- failureThreshold: 90: Tolerates up to 90 consecutive failures. **Total grace period = 90 × 20 = 1,800 sec (30 min).**
- Kubernetes waits up to 30 mins for the service to become healthy before marking it failed.
- It prevents premature pod restarts during heavy startup which is critical for large-scale deployments with many devices.
- It gives ample time for InventoryDeviceSweep to complete even in large-scale deployments.

Tenant VRF Update Fails

Updating a tenant VRF using the **efa tenant vrf update** command fails with an error message indicating that when a VRF is in the `vrf-device-created` state, you cannot add or delete a centralized router.

```
efa:extreme)extreme@efa:~$ efa tenant vrf update --tenant "<tenant-name>" --name "<vrf-name>" --operation=centralized-router-add --centralized-router x.x.x.x
Vrf updation Failed:
Error : Operation "centralized-router-add" is not allowed for Vrf in "vrf-device-created" state
```

Resolution

If you have an existing centralized router and want to add another one, complete the following steps:

1. Delete the BGP peer-group associated with the BGP (which is linked to the VRF being updated) for the existing centralized router.

```
efa tenant service bgp peer-group delete --tenant "<tenant-name>" --name "<pg-name>" (--force)
```

2. Delete the BGP peer associated with the VRF for the existing centralized router.

```
efa tenant service bgp peer delete --tenant "<tenant-name>" --name "<bgp-peer-instance-name>" (--force)
```

3. Delete the EPG associated with the VRF being updated.

```
efa tenant epg delete --tenant "<tenant-name>" --name "<epg-name>" (--force)
```

4. Delete the VRF being updated.

```
efa tenant vrf delete --tenant "<tenant-name>" --name "<vrf-name>"
```

5. Recreate the tenant VRF, EPG, BGP peer-group, and BGP peers.

```
efa tenant vrf create --tenant "<tenant-name>" --name "<vrf-name>" --routing-type"centralized" --layer3-extension-enable false --centralized-router "x.x.x.x,y.y.y.y" --local-asn <asn> --rh-ecmp-enable --rh-max-path 64 --rt-typeimport --rt 101:101 --rt-typeexport --rt 101:101 --max-path 64 --redistribute connected --graceful-restart-enable
efa tenant epg create --tenant "<tenant-name>" --name "<epg-name>" --type extension --switchport-mode trunk --single-homed-bfd-session-type auto --po <po-name> --vrf <vrf-name> --ctag-range <vlan-number> --anycast-ip <vlan-number>:z.z.z.z/24 --local-
```

```
ip <vlan-number>,a.a.a.a:b.b.b.b/24 --local-ip <vlan-number>,c.c.c.c:d.d.d.d/24 --ctag-
description "<description"
efa tenant service bgp peer-group create --tenant "<tenant-name>" --name
"<pg-name>" --pg-name <device-ip:peer-group-name> --pg-asn <device-ip,peer-group-
name:remote-asn> --pg-bfd-enable <device-ip,peer-group-name:bfd-enable(true/false)> --
pg-bfd <device-ip,peer-group-name:bfd-interval,bfd-min-rx,bfd-multiplier> --pg-update-
source-ip <device-ip,peer-group-name:update-source-ip>
efa tenant service bgp peer create --tenant "<tenant-name>" --name "<bgp-peer-
instance>" --ipv4-uc-dyn-nbr <device-ip,vrf-name:ipv4-listen-range,peer-group-name>
```



Note

- If there was no existing centralized router configuration, start from step 3.
- You can add or delete the centralized router if the VRF state is `vrf-create`.

Unable to Delete an EPG with Comma in Name

When attempting to delete an Endpoint Group (EPG) with a comma (",") in its name, the system interprets the comma as a delimiter and treats it as multiple EPGs. This results in an error, as deletion of multiple EPGs is not supported.

```
efa tenant epg delete --tenant VR6-Tnt --name "epg52,56"
```

```
Error : deletion of multiple EPG is not supported. Only one EPG can be deleted at a time
```

Resolution

Delete the EPG using REST API.



Note

- The following is an example query to delete an EPG using VSCode with the 'REST Client' extension:

```
@hostname = https://xxx.xxx.xxx.xxx # Add EFA IP here
@access-token = # Paste bearer token here
```

- The following is an example query to get bearer access token:

```
POST {{hostname}}/v1/auth/token/access-token
Content-Type: application/json

{
  "username": "extreme",
  "password": "password"
}
```

From the output, copy the access token and paste it into the @access-token variable.

- The following is an example query to delete an endpoint group:

Change the tenant and EPG name as needed.

```
DELETE {{hostname}}/v1/tenant/endpointgroup
Content-Type: application/json
Authorization: Bearer {{access-token}}

{
  "tenant_name": "example-tenant",
  "endpoint-group-list": [
    {
      "name": "example-epg"
    }
  ]
}
```

You can use the API client of your choice.

VLAN Configuration Fails When Adding Port or PO to EPG

When attempting to add a port channel (PO) to an Endpoint Group (EPG), the following error occurs:

```
efa tenant epg update --name VLAN0100_Tagged --po PO_2 --operation port-group-add --
tenant Tnt1
Device: 10.10.10.29 Error: Configuring switchport mode trunk on 2 failed due to netconf
rpc [error] '% Error: One or more VLANs are not configured.
```

The error occurs because the SLX device lacks the required VLAN configuration. This causes EFA to fail when trying to add VLAN to the port or PO.

Resolution

Manually add the VLAN to the SLX device before running the EPG command.



Note

- As a best practice, consider upgrading to a supported release. If the issue persists, contact TAC for further assistance.
- To assign VLANs to ports (create a VLAN), complete the following steps:
 1. Create a tenant using the following command:

```
efa tenant create --name <tenant-name> --vlan-range <vlan-id-range> --port
<device-ip>[<slot/port>]
```

2. Create an Endpoint Group (EPG) associated with the tenant using the following command:

```
efa tenant epg create --name <epg-name> --tenant <tenant-name> --port
<device-ip>[<slot/port>] --switchport-mode <type of switchport> --ctag-
range <vlan-id-range>
```

This will result in the following configuration on the SLX device:

```
interface Ethernet <slot/port>]
  switchport
  switchport mode trunk
  switchport trunk allowed vlan add <vlan-id>
  no switchport trunk tag native-vlan
  no shutdown
```

VRF Already Exists on Device Error

When attempting to create a new VRF on an SLX device managed by XCO Fabric Skill, the following error occurs:

```
(efa:extreme)extreme@PD2-S01-tpvm:~$ efa tenant vrf create --tenant "VR7-Test" --name
"testVR"
Error : VRF with name testVR already exist on device, use different VRF name
```

This error occurs when XCO checks the device for existing configured VRFs while adding a new VRF, and finds that the target VRF name ("testVR") is already in use. Alternatively, if the VRF was previously deleted, but XCO's inventory device has not been updated, the same error may occur.

```
PD2_BL1# show run vrf | inc testVR
vrf testVR
PD2_BL1# conf
Entering configuration mode terminal
PD2_BL1(config)# no vrf testVR

(efa:extreme)extreme@PD2-S01-tpvm:~$ efa tenant vrf create --tenant "VR7-Test" --name
"testVR"

Error : VRF with name testVR already exist on device, use different VRF name
```

Resolution

1. Manually delete the VRF if it exists on the target device.

- a. Enter configuration mode.

```
PD2_BL1# conf
Entering configuration mode terminal
```

- b. Delete the VRF.

```
PD2_BL1(config)# no vrf testVR
```

- c. Exit configuration mode.

```
PD2_BL1(config)# end
PD2_BL1#
```

2. Update the inventory device via XCO.

```
(efa:extreme)extreme@PD2-S01-tpvm:~$ efa inventory device update --ip 10.53.4.44-45
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | IP Address | Host Name | Model | Chassis Name
| Firmware | ASN | Role | Fabric | Status | Reason |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 13 | 10.53.4.44 | PD2_BL1 | 4004 | SLX9740-40C | 20.4.2b
| 4200003000 | BorderLeaf | pd2_cnis | Success | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 15 | 10.53.4.45 | PD2_BL2 | 4004 | SLX9740-40C |
20.4.2b | 4200003000 | BorderLeaf | pd2_cnis | Success | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Device Details
--- Time Elapsed: 15.988729903s ---
```

3. Create the VRF in XCO.

```
(efa:extreme)extreme@PD2-S01-tpvm:~$ efa tenant vrf create --tenant "VR7-Test" --name
"testVR"
```

```
Vrf created successfully.
--- Time Elapsed: 443.152631ms ---
```




Note

- The VRF will not be configured onto the device until an EPG is associated with it.
- Verify the VRF creation on the SLX device using the **show run vrf** command.

```
PD2_BL1# show run vrf | inc testVR
PD2_BL1#
PD2_BL2# show run vrf | inc testVR
PD2_BL2#
```

- Verify the VRF details in XCO using the **efa tenant vrf show** command.

```
(efa:extreme)extreme@PD2-S01-tpvm:~$ efa tenant vrf show --name testVR --tenant VR7-Test
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Name | Tenant | Routing Type | Centralized | Enable L3 | Redistribute |
| Max | Local | Enable | State | Dev State | App State |
| Path | Asn | GR | Routers | Extension | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| testVR | VR7-Test | distributed | | true | connected |
| 8 | | false | vrf-created | not-provisioned | cfg-ready |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Vrf Details
--- Time Elapsed: 326.533647ms ---
```



Troubleshooting Upgrade and Migration

- [Active Node is Down after EFA or XCO Upgrade](#) on page 111
- [EFA or XCO Upgrade Fails on Ubuntu Server](#) on page 112
- [EFA or XCO Upgrade to 3.2.x Fails](#) on page 112
- [Incorrect EFA or XCO Version Displayed After Upgrade](#) on page 113
- [Kubernetes POD CoreDNS Stuck in CrashLoopBackOff Post Upgrade](#) on page 113
- [Login to EFA or XCO Fails after Upgrade](#) on page 114
- [SLX Upgrade Failure and Inability to Restore](#) on page 116
- [TPVM Incremental Upgrade Fails](#) on page 117
- [TPVM Upgrade on the Standby Node Fails During a Data Fabric Upgrade](#) on page 117
- [User Accounts Deleted Post TPVM Upgrade](#) on page 118
- [XCO Firmware Upgrade Fails](#) on page 118
- [XCO or EFA Lost Connection to Devices After Upgrade](#) on page 118
- [EFA Status Remains Down After Upgrade](#) on page 120

Learn about troubleshooting upgrade and migration issues.

Active Node is Down after EFA or XCO Upgrade

After upgrading EFA or XCO, the active node status is displayed as "down".

```
(efa:extreme)extreme@b144-s1:~$ efa status
+-----+-----+-----+-----+
| Node Name | Role   | Status | IP           |
+-----+-----+-----+-----+
| b144-s1   | active | down   | 21.144.144.2 |
+-----+-----+-----+-----+
| b144-s2   | standby | up     | 21.144.144.3 |
+-----+-----+-----+-----+
```

This issue occurs because EFA or XCO is not properly handling the deletion of a kube-system pod.

Resolution

Run the following command:

```
TPVM# k3s kubectl -n kube-system delete ep/kube-controller-manager
```

EFA or XCO Upgrade Fails on Ubuntu Server

When upgrading EFA or XCO on a standalone Ubuntu server, the upgrade process encounters a failure during the `Installing mariadb server` stage.

The failure occurred due to NIS/YP configuration on the server, leading to a `whiptail` prompt, which is not supported in the text-based upgrades.

```
whiptail --backtitle "Package configuration" --title "Configuring mariadb-server-10.6" --nocancel --msgbox "Important note for NIS/YP users..."
```

Resolution

To allow the installation to continue and complete successfully, complete the following steps:

1. Kill the hanging `whiptail` process to allow the installation to continue and complete successfully.
2. Remove the `debconf` configurations for `mariadb-server` (or set the `DEBIAN_FRONTEND` environment variable to `noninteractive`) to prevent it from applying its settings during the upgrade.
3. Run the following command to purge MySQL or MariaDB packages and ensure a clean installation environment for the upgrade:

```
DEBIAN_FRONTEND=noninteractive apt-get remove --purge -y -o Dpkg::Options::="--force-confdef" -o Dpkg::Options::="--force-confold" mysql-server mysql-client mysql-common mysql-server-core-* mysql-client-core-*; apt-get -y autoremove; apt-get clean all;
```

EFA or XCO Upgrade to 3.2.x Fails

The upgrade from EFA 2.7.x to XCO 3.2.x is unsuccessful.

Despite attempting the upgrade with the specified options (EFA 3.2.x build GA, ipv4 IP Stack, and Fabric Automation suites), the process fails.

The installation fails while checking if the `monitor` service is running, prompting the user to unwind the partial installation and collect a `supportsave`.

```
Verifying if monitor service is running...
Failed.
Installation failed. Do you want to unwind the partial install? (yes/no)
yes
Do you want to collect a supportsave? (yes/no)
yes
{"@time":"2022-07-26T11:08:44.172987 CEST","App":"common","msg":"reliable_k3s kubectl
apply -f yaml/coredns.yaml"}
common.sh: line 15280: k3s: command not found
```

The logs indicate repeated attempts and possible interruptions during the upgrade process, resulting in errors during removal of the previous version and multiple reinstall attempts. A critical error occurs when running the `kubectl apply` command, as the `k3s` command is not found (line 15280 in `common.sh`).

Resolution

1. Perform a fresh install of Ubuntu OS.
2. Update the "PATH" variable setting globally to ensure proper functioning of services.

Check whether `/usr/local/bin` is added in the following secure path:

```
Defaults
secure_path="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/
bin:/snap/bin"
```

Alternatively, add `/usr/local/bin` and then redeploy.



Note

- If a backup tar from efa-2.7.x is available, perform a fresh install of XCO 3.2.x on a new Ubuntu OS, followed by a database restore.
- After restoring the database, ensure that device certificates are updated.

Incorrect EFA or XCO Version Displayed After Upgrade

When you remain connected to TPVM during an EFA or XCO upgrade, the `efa version` command may still display the older version instead of the updated one.

Resolution

To confirm the EFA or XCO version, log out of TPVM and EFA or XCO, and then log back in.

Kubernetes POD CoreDNS Stuck in CrashLoopBackOff Post Upgrade

After upgrading from EFA 2.x.x to XCO 3.x.x, the Kubernetes POD `coredns-6f5nw` is stuck in a **CrashLoopBackOff** state.

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE	NOMINATED
NODE	READINESS GATES						
pod/coredns-6f5nw	0/1	CrashLoopBackOff	13 (4m7s ago)	31m	10.42.0.2	efa	

This issue occurs because the default DNS resolution on Ubuntu 18 or later uses `systemd-resolved`, which provides a local DNS stub listener on `127.0.0.53`.

The `/etc/resolv.conf` file is managed by `systemd-resolved` and contains a single `nameserver` entry (`127.0.0.53`).

```
# This file is managed by man:systemd-resolved(8). Do not edit.
nameserver 127.0.0.53
options edns0
```

The `systemd-resolve` listens on this IP.

```
sudo ss -tulpan | grep "127.0.0.53"
udp UNCONN 0 0 127.0.0.53%lo:53 0.0.0.0:* users:(("systemd-
resolve",pid=685,fd=12))
```

```
tcp LISTEN 0 128 127.0.0.53%lo:53 0.0.0.0:* users:(("systemd-
resolve",pid=685,fd=13))
```

As a result, CoreDNS is configured to forward DNS queries to this address. The following is an example of Corefile:

```
.:53 {
  errors
  health
  ready
  kubernetes cluster.local in-addr.arpa ip6.arpa {
    pods insecure
    fallthrough in-addr.arpa ip6.arpa
  }
  hosts /etc/coredns/NodeHosts {
    ttl 60
    reload 15s
    fallthrough
  }
  prometheus :9153
  forward . /etc/resolv.conf
  cache 30
  loop
  reload
  loadbalance
}
```

Resolution

1. Ensure that a valid nameserver is configured on the host by running the **sudo systemd-resolve --status** command.

The following is an example output:

```
Link 2 (ens160)
  Current Scopes: none
  LLMNR setting: yes
  MulticastDNS setting: no
  DNSSEC setting: no
  DNSSEC supported: no
```

2. Upgrade XCO to version 3.x.x or later. If no DNS server is configured on the server, CoreDNS will crash in **CrashLoopBackOff**.
3. Disable the access of XCO services to the host DNS nameserver by running the **update-dns.sh** script with the `--dns-action disallow` option.

The script location varies depending on the environment (TPVM based deployment or server based deployment).

```
sudo <location_of_script>/update-dns.sh --dns-action disallow
```

- On TPVM


```
/apps/efa/update-dns.sh --dns-action disallow
```
- On a server


```
/opt/efa/update-dns.sh --dns-action disallow
```

Login to EFA or XCO Fails after Upgrade

After upgrading EFA or XCO, login attempts are unsuccessful.

The `efact1 status` command output shows that the `gosystem-service` POD is stuck in a `CrashLoopBackOff` state, and several PODs are either not ready or initializing.

```
EFA Application Status:
NAME                                READY   STATUS             RESTARTS   AGE   IP              NODE
NOMINATED NODE  READINESS GATES
pod/gosnmp-service-mrt6d            0/1    Init:0/1          0           16h   10.158.129.69   efa
<none>                               <none>
pod/gorbac-service-xkl6b            0/1    Init:0/2          0           16h   10.42.194.122   efa
<none>                               <none>
pod/gotentant-service-pcxv2         0/1    Init:0/3          0           16h   10.42.194.123   efa
<none>                               <none>
pod/efa-api-docs-bzcrb              1/1    Running           0           16h   10.42.194.124   efa
<none>                               <none>
pod/goraslog-service-xl8lp          0/1    Init:0/3          0           16h   10.158.129.69   efa
<none>                               <none>
pod/gohyperv-service-8xsb7          0/1    Init:0/2          0           16h   10.42.194.66    efa
<none>                               <none>
pod/govcenter-service-4x4qt         0/1    Init:0/2          0           16h   10.42.194.127   efa
<none>                               <none>
pod/goinventory-service-zhp9d       0/1    Init:0/2          0           16h   10.42.194.67    efa
<none>                               <none>
pod/gofabric-service-x4zbd          0/1    Init:0/2          0           16h   10.42.194.65    efa
<none>                               <none>
pod/gonotification-service-pttss    0/1    Init:0/2          0           16h   10.158.129.69   efa
<none>                               <none>
pod/goauth-service-dmgkk            0/1    Init:0/2          0           16h   10.42.194.73    efa
<none>                               <none>
pod/goopenstack-service-qf428       0/1    Init:0/1          0           16h   10.42.194.68    efa
<none>                               <none>
pod/rabbitmq-jpjml                  1/1    Running           0           16h   10.42.194.69    efa
<none>                               <none>
pod/gopolicy-service-6ltrq          0/1    Init:2/3          0           16h   10.42.194.125   efa
<none>                               <none>
pod/gosystem-service-x44n6          0/1    CrashLoopBackOff  194          16h   10.42.194.126   efa
<none>                               <none>
```

The issue is caused by a failed MariaDB service startup, which was previously shut down normally.

database/error.log:

```
2023-04-15 2:22:13 0 [Note] /usr/sbin/mysqld (initiated by: unknown): Normal shutdown
2023-04-15 2:22:13 0 [Note] Event Scheduler: Purging the queue. 0 events
2023-04-15 2:22:13 0 [Note] InnoDB: FTS optimize thread exiting.
2023-04-15 2:22:13 0 [Note] InnoDB: Starting shutdown...
2023-04-15 2:22:13 0 [Note] InnoDB: Dumping buffer pool(s) to /opt/efadata/mysql/ib_buffer_pool
2023-04-15 2:22:13 0 [Note] InnoDB: Instance 0, restricted to 4013 pages due to
innodb_buf_pool_dump_pct=25
2023-04-15 2:22:13 0 [Note] InnoDB: Buffer pool(s) dump completed at 230415 2:22:13
2023-04-15 2:22:15 0 [Note] InnoDB: Shutdown completed; log sequence number 317648808; transaction id
965556
2023-04-15 2:22:15 0 [Note] InnoDB: Removed temporary tablespace data file: "ibtmp1"
2023-04-15 2:22:15 0 [Note] /usr/sbin/mysqld: Shutdown complete
```

database/systemctl_10.158.129.69.log:

```
* mariadb.service
  Loaded: masked (/dev/null; bad)
  Active: inactive (dead) since Sat 2023-04-15 02:22:15 CEST; 3 days ago
Main PID: 992 (code=exited, status=0/SUCCESS)
  Status: "MariaDB server is down"

Mar 17 09:32:39 seliius22899 systemd[1]: Starting MariaDB 10.4.17 database server...
```

```

---
Mar 17 09:32:43 seliius22899 /etc/mysql/debian-start[2088]: FATAL ERROR: Upgrade failed
Apr 15 02:22:13 seliius22899 systemd[1]: Stopping MariaDB 10.4.17 database server...
Apr 15 02:22:15 seliius22899 systemd[1]: Stopped MariaDB 10.4.17 database server.

```

Resolution

Manually start the MariaDB service using the `sudo systemctl start mariadb` command.

SLX Upgrade Failure and Inability to Restore

The upgrade of an SLX 8720 switch fails, and resulted in an inability to restore the previous version. The following symptoms are observed:

1. A non-Clos cluster with two SLX 8720 switches:
 - SLX 8720 SLXLB
 - SLX 8720 SLXLA
2. Upgrade of SLXLB failed with `Device Reload Failed` error, and it is down.
3. Switch restoration to the previous version using the restore command (**SLX# firmware restore**) fails due to invalid credentials.

After the upgrade failure, credentials stopped working, causing the restore to fail with an `Invalid credentials` error.

```

(efa:extreme)extreme@tpvmslx1a:~$ efa inventory device compare --ip x.x.x.x
Error: Failed to establish connection with device (x.x.x.x), Error: Invalid credentials for device
x.x.x.x
Compare the device configuration with the configuration details saved in the application.
Usage:
  efa inventory device compare [flags]
Flags:
  --ip string  IP Address of the device to be compared
  --- Time Elapsed: 10.594416161s ---
Error: Invalid credentials for device x.x.x.x

```

4. The SLX upgrade is also unsuccessful when using the default login credentials.

The issue occurs because EFA or XCO was not installed in a multi-node environment.

Resolution

To recover and restore the system, complete the following steps:

1. Un-deploy EFA or XCO from SLXLB.
2. Redeploy EFA or XCO on both SLXA (master) and SLXLB (standby) with secondary management interface.

EFA or XCO deployment is successful.

3. Restore EFA or XCO system from the backup file.
4. Perform EFA or XCO DRC on both SLXA and SLXLB.

**Note**

EFA or XCO was initially installed only on SLXB, but reinstallation on multiple nodes is necessary for recovery.

TPVM Incremental Upgrade Fails

The incremental upgrade of TPVM from version 4.5.1-0 to 4.5.2-1 is unsuccessful.

Resolution

The TPVM incremental upgrade is supported only from EFA version 3.0.0 and later. Prior to EFA 3.0.0, you must perform a full upgrade, which involves removing the old TPVM, creating a new TPVM with the updated version, and then deploying EFA on it. This process can take around 50 minutes and requires individual upgrade for each TPVM node, starting with the standby node followed by the active node. This process included a TPVM reboot, resulting in EFA service downtime.

However, if you meet the upgrade requirements, incremental upgrades offer a faster and more efficient process, taking approximately 3 minutes, and can be applied to both active and standby TPVM nodes, as well as to single-node deployments, provided the SLX version is greater than 20.4.1.

TPVM Upgrade on the Standby Node Fails During a Data Fabric Upgrade

During a data fabric upgrade, the TPVM upgrade to version 4.6.7 on the standby node fails due to an unexpected inventory service restart. Prior to this, EFA or XCO had been upgraded to 3.4.0, and the SLX devices were successfully upgraded to 20.5.3.

The following symptoms are observed:

1. Connection to EFA or XCO is lost during the upgrade process.
2. Inventory service restarts unexpectedly on the active node during TPVM full upgrade.
3. Both TPVMs showed down status initially, then active TPVM recovered, but standby TPVM remains down after the upgrade attempt.
4. Old TPVM software version (4.5.14) still shown in EFA inventory .

The issue occurred when the active node lost connectivity, unable to reach either the peer node or the default gateway. This led to the High Availability (HA) system activating its self-fencing mechanism, which caused the inventory-service to restart and, in turn, disrupted the TPVM upgrade process on the standby node.

Resolution

The root cause of the issue is unclear, but it's related to network connectivity issue.

- Auto-reboot parameter was forcefully set to false, disabling auto-reboot functionality
- Further investigation is needed to determine why the active node couldn't ping the default gateway during the upgrade process.
- Manual intervention may be required to complete TPVM upgrade.

User Accounts Deleted Post TPVM Upgrade

After upgrading to any version of TPVM, all custom-created user accounts, including usernames and passwords, are permanently deleted.

This issue occurs because the upgrade process involves a full uninstallation of the previous TPVM version before installing the new one, leading to the loss of custom-configured accounts.

Resolution

1. To regain access, perform an incremental TPVM upgrade.



Note

Incremental TPVM upgrades are supported starting from EFA 3.0.0.

2. Recreate all deleted user accounts, including usernames and passwords.

XCO Firmware Upgrade Fails

Unable to upgrade SLX9250 from SLXOS 20.5.2a to 20.5.3a using XCO 3.4.0, resulting in the following error message:

```
Cannot find firmware. The server is inaccessible or firmware path is invalid. Please verify that the server name or IP address, user/password, and firmware path are correct.
```

The issue occurred due to incorrect FTP or SCP server and firmware path configuration.

Resolution

Ensure that the FTP or SCP server IP address and directory path to the firmware are configured properly.

XCO or EFA Lost Connection to Devices After Upgrade

Use this topic to learn about troubleshooting XCO or EFA connection failures to devices after upgrade.

Scenario 1: Connection Loss and Certificate Error

After upgrade, XCO loses connection to devices, resulting in the following error log:

```
"time":"2023-04-14T15:39:34.75238 CEST","level":"error","msg":"HTTPS adapter login error: Get \"https://x.x.x.x/rest\": x509: certificate signed by unknown authority (possibly because of \"crypto/rsa: verification error\" while trying to verify candidate authority certificate \"EFA Intermediate CA\")"
```

This issue occurs when running the following commands:

1. **efa certificate device install -ip <slx-ip>**
2. **efa inventory device update --ip <slx-ip>**

Resolution

As a best practice, reinstall certificates as a recovery procedure.



Note

Reinstallation of certificates is crucial as part of the backup restoration process. Including this step in testing procedures ensures maintenance mode functions correctly upon reboot.

Scenario 2: Fallback Communication Failure

After upgrading from EFA 3.1.x to XCO 3.2.x, EFA or XCO fails to communicate with most devices in the fabric due to missing certificates. The SLX certificate files had disappeared from EFA directory (/apps/efadata/certs/) for the affected switches. Although the certificates were still valid on the switches, EFA had removed them, causing communication failures. It was observed that EFA 3.1.x had a secondary (fallback) communication method that was missing in EFA 3.2.x, leading to communication issues post upgrade.

When attempting to run a CLI command on the affected device, the following error was encountered:

```
efa inventory device execute-cli --ip 192.168.246.32 --command "show version"
Execute CLI[failed]
-----
IP Address      Host Name      Fabric      Command      Status      Output
Reason
-----
192.168.246.32  RNOC3000LEAF1B35  south-5gc-fabric  show version  Failed      Device 192.168.246.32
not reachable. Please retry after verifying the inputs and connectivity issues.
```

You can verify the certificates on switches using the **show crypto ca certificates** command.

Resolution

To resolve the communication issue, reinstall the certificates on the affected devices using the **efa certificate device install** command.

EFA Status Remains Down After Upgrade

During an end-to-end upgrade procedure, the EFA status remains down after upgrading XCO from version 3.4.x to later.

After completing the XCO upgrade, the following symptoms are observed:

- EFA status shows as "down".
- EFA version shows as 3.5.0, but status remains down.

```
extreme@dc276-slx-11a-tpvm:~$ efa version
Version: 3.5.0
Build: GA
Time Stamp: 24-04-08:13:23:50
Mode: Secure
Deployment Type: multi-node
Deployment Platform: TPVM
Deployment Suite: Fabric Automation
Deployment IP Mode: ipv4
Virtual IP: 192.168.246.2
Node IPs: 192.168.246.3,192.168.246.4
Ping Target Enabled: yes
HA Health Check IP(s): 192.168.246.1
- Time Elapsed: 9.135809ms -
(efa:no-auth)extreme@dc276-slx-11a-tpvm:~$ efa status
-----+
Node Name           Role      Status   IP
-----+
dc276-slx-11a-tpvm  active   down    192.168.246.3
-----+
dc276-slx-11b-tpvm  standby  down    192.168.246.4
-----+
- Time Elapsed: 14.311591578s -
(efa:no-auth)extreme@dc276-slx-11a-tpvm:~$ efacl status
Node: efa
(efa:no-auth)extreme@dc276-slx-11a-tpvm:~$
```

There are two possible causes for the issue:

1. The default gateway settings are not consistent between TPVM and SLXs.
2. External services like NTP and DNS cannot access the default gateway.

Resolution

1. Ensure that the default gateway configuration is aligned between TPVM and SLXs.
2. Ensure that the default gateway is consistent and reachable by external entities like NTP and DNS.

The following is an example configuration of a proper default gateway settings:

```
DC276-SLX-L1A# show interface Management
interface Management 0
line-speed actual "1000baseT, Duplex: Full"
oper-status up
rme-info "mgmt2 mgmt1(AP)"
ip address "static 192.168.246.11/24"
ip gateway-address 192.168.246.1
```

```
DC276-SLX-L1A# config
Entering configuration mode terminal
DC276-SLX-L1A(config)# tpvm
DC276-SLX-L1A(config-tpvm-TPVM)# interface management ip 192.168.246.3/24 gw
192.168.246.11
DC276-SLX-L1A(config-tpvm-TPVM)# end

DC276-SLX-L1B# config
Entering configuration mode terminal
DC276-SLX-L1B(config)# tpvm
DC276-SLX-L1B(config-tpvm-TPVM)# interface management ip 192.168.246.4/24 gw
192.168.246.11
DC276-SLX-L1B(config-tpvm-TPVM)# end

DC276-SLX-L1A# efa deploy non-interactive multi-node package /efaboot/efa-3.5.0.tar.gz
peer-node 192.168.246.4 vip4 192.168.246.2 ping-target 192.168.246.1
```