Extreme Fabric Automation Release Notes

Version 2.5.0

9037132-00 Rev AA
July 2021

Legal Notice

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

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

Trademarks

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

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

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

Open Source Declarations

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

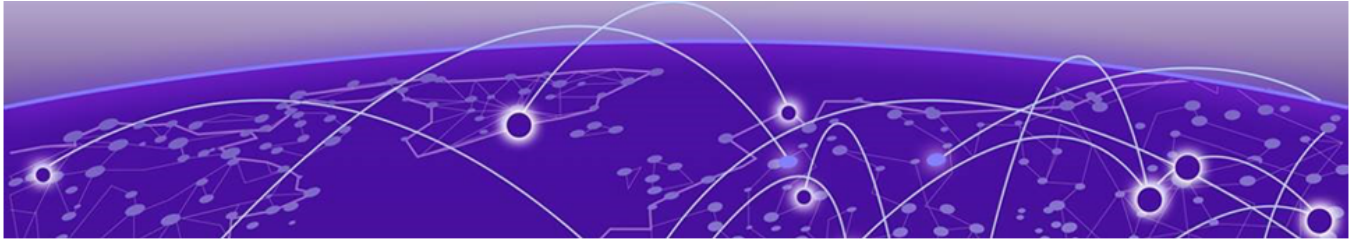

Table of Contents

Preface

Read the following topics to learn about:

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

Text Conventions

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

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

Table 1: Notes and warnings

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

Table 2: Text

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

Table 3: Command syntax

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

Documentation and Training

Find Extreme Networks product information at the following locations:

- Current Product Documentation
- Release Notes
- Hardware and software compatibility for Extreme Networks products
- Extreme Optics Compatibility
- Other resources such as white papers, data sheets, and case studies

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

Getting Help

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

Extreme Portal

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

The Hub

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

Call GTAC

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

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

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

Subscribe to Product Announcements

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

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

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

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, 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 in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

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

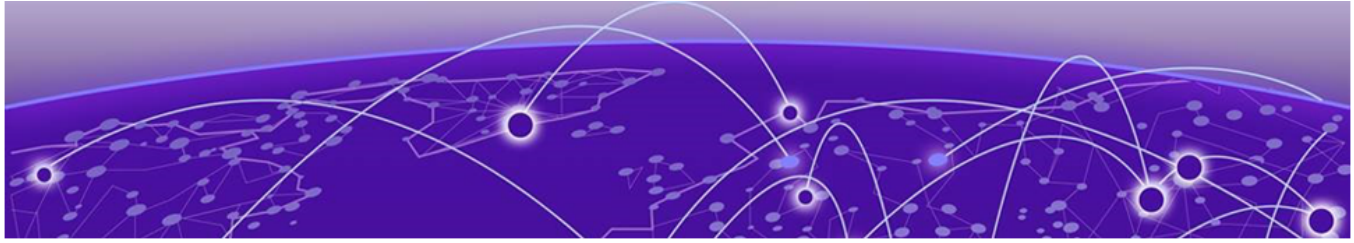

Release Notes

New in this Release

Extreme Fabric Automation 2.5.0 provides the following features and improvements.

Table 4: Features and improvements

Feature	Description
Suppressing or enabling ARP and neighbor discovery	You can use new parameters for the **`efa tenant epg create`** command to suppress or enable Address Resolution Protocol (ARP) or neighbor discovery. • **`--suppress nd`**: Set to false to enable neighbor discovery. Set to true to suppress neighbor discovery. • **`--suppress arp`**: Set to false to enable ARP. Set to true to suppress ARP. For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.
Software Bidirectional Forwarding Detection (BFD) sessions on the Cluster Edge Port (CEP) on SLX 9150 and SLX 9250 devices	The **`efa tenant epg create`** command provides the **`single-homed-bfd-session-type`** parameter, with which you can set the BFD session type to software, hardware, or auto. For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.

Table 4: Features and improvements (continued)

Feature	Description
Shared VRF and endpoint groups for a shared tenant	You can use the following commands to configure shared VRF and shared endpoint groups. • **efa tenant vrf create** • **efa tenant epg create** For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.
Local ASNs per tenant VRF	You can use the new **--local-asn** parameter for the **efa tenant vrf create** and **efa tenant vrf update** commands to configure the local ASN during creation and update activities. For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.
SNMP, NTP, and MTU configuration persists after upgrade and is then managed by EFA	After upgrade, use the EFA CLI for changes to these configurations. For more information, see the *Extreme Fabric Automation Administration Guide, 2.5.0*.
NTP configuration on an SLX device persists in the EFA database	You can configure NTP at the device and fabric levels using the **efa inventory device ntp server create** command. For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.
Redundant Management Ethernet configuration on an SLX interface persists in the EFA database	For SLX 9150, SLX 9250, and SLX 9740, use the new **efa inventory device interface redundant-management** command to configure RME. For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.
Enabling or disabling telemetry, specifically the operational state for all modules including BGP, interface, and platform	Use the **efa inventory device telemetry** command with the **--enable bool** parameter. For more information, see the *Extreme Fabric Automation Command Reference, 2.5.0*.
TPVM upgrade while EFA is running	You can upgrade a TPVM image where TPVM is installed and running with an EFA instance. For more information, see the *Extreme Fabric Automation Deployment Guide, 2.5.0*.
Extra routes in Neutron	API with basic ECMP (non-resilient hashing) on the OpenStack Layer 3 plugin. For more information, see the *Extreme Fabric Automation OpenStack Integration Guide, 2.5.0*.

Table 4: Features and improvements (continued)

Feature	Description
OpenStack journaling	Enhancements have been added to Journaling: • Detection of EFA reachability and early failure of OpenStack commands • Limit of journal retries if the errors are non-recoverable until efa-sync. For more information, see the *Extreme Fabric Automation OpenStack Integration Guide, 2.5.0*.
Security hardening	A new guide describes such topics as user authentication, MD5 authentication, and security hardening. For more information, see the *Extreme Fabric Automation Security Guide, 2.5.0*.

CLI Commands

There are new, modified, and deprecated commands in this release.

New commands

- **efa inventory device interface list-breakout**
- **efa inventory device interface redundant-management**
- **efa inventory device interface set-fec**
- **efa inventory device interface set-link-error-disable**
- **efa inventory device interface unset-fec**
- **efa inventory device interface unset-link-error-disable**
- **efa inventory device ntp disable-server**
- **efa inventory device ntp server create**
- **efa inventory device ntp server delete**
- **efa inventory device ntp server list**
- **efa inventory device snmp community create**
- **efa inventory device snmp community delete**
- **efa inventory device snmp community list**
- **efa inventory device snmp host create**
- **efa inventory device snmp host delete**
- **efa inventory device snmp host list**
- **efa inventory device snmp user create**
- **efa inventory device snmp user delete**
- **efa inventory device snmp user list**
- **efa inventory device timezone**
- **efa tenant service bgp peer operational show**

- **efa tenant vrf error show**
- **openstack network efa-bl-pair-map create**

Modified commands

- **efa fabric setting update**
- **efa inventory device add-bulk**
- **efa inventory device firmware-download execute**
- **efa inventory device firmware-download prepare add**
- **efa inventory device firmware-download prepare remove**
- **efa system backup**
- **efa system settings update**
- **efa inventory device interface set-admin-state**
- **efa inventory device interface set-breakout**
- **efa inventory device interface set-mtu**
- **efa inventory device interface set-speed**
- **efa inventory device interface unset-breakout**
- **efa tenant epg create**
- **efa tenant service bgp peer create**
- **efa tenant service bgp peer update**

Deprecated commands

efa inventory device system set-mtu

Supported Platforms and Deployment Models

Support includes bare metal, OVA, and TPVM deployment models, supported TPVM versions, supported SLX-OS software versions, and supported SLX devices.

Table 5: Bare Metal Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
EFA 2.3.x, 2.4.x, and 2.5.x	External server (bare metal)	More than 24	Yes	16.04 and 18.04	• CPU: 4 cores • Storage: 50 GB • RAM: 8 GB

Table 6: OVA Deployment Models

EFA Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
2.3.x, 2.4.x, 2.5.x (Secure mode)	External server (OVA)	More than 24	Yes	18.04	• CPU: 4 cores • Storage: 50 GB • RAM: 8 GB

Table 7: TPVM Deployment Models

EFA Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
2.3.x	• SLX 9150 • SLX 9250 • SLX 9740	Up to 24	Yes	18.04	20.2.2a
2.4.x	• SLX 9150 • SLX 9250 • SLX 9740	Up to 24	Yes	18.04	20.2.2b
2.5.x	• SLX 9150 • SLX 9250 • SLX 9740	Up to 24	Yes	18.04	20.2.3.f

Table 8: TPVM Software Support

TPVM Version	SLX-OS 20.2.3d/e/f	SLX-OS 20.3.2	SLX-OS 20.3.2a	Ubuntu Version	EFA Version
4.2.2	Yes	No	No	18.04	2.3.x
4.2.4	Yes	No	No	18.04	2.4.x
4.2.5	No	Yes	Yes	18.04	2.4.x, 2.5.x

Note
The seamless TPVM upgrade feature is not available in SLX 20.2.3f.

Table 9: IP Fabric Topology Matrix

Device	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9150	20.1.x, 20.2.x, 20.3.x	✔				✔
SLX 9250	20.1.x, 20.2.x, 20.3.x	✔	✔			✔
SLX 9540	20.1.x, 20.2.x, 20.3.x	✔			✔	
SLX 9640	20.1.x, 20.2.x, 20.3.x				✔	
SLX 9740	20.2.x, 20.3.x		✔	✔	✔	✔

Table 10: EFA, Neutron, and SLX-OS Compatibility

EFA Version	Neutron Version	SLX-OS Version
2.3.0	2.3.0_19	20.1.2d
2.3.1, 2.3.2	2.3.1_02	20.1.2e, 20.2.2a
2.4.0, 2.4.1	3.0.0-23	20.2.3, 20.2.3a/b/c
2.4.2, 2.4.3, 2.4.4, 2.4.6	3.0.1-07	20.2.3d/e/f
2.5.0	3.1.0-15	20.3.2a

EFA Installation and Upgrade

You can install and upgrade EFA in single-node environments, in multi-node high-availability environments, and on a TPVM. You can also upgrade from a single-node to a multi-node deployment.

For complete information about installation and upgrade scenarios, see the *Extreme Fabric Automation Deployment Guide, 2.5.0*.

Defects Closed with Code Changes

The following defects, which were previously disclosed as open, were resolved in Extreme Fabric Automation 2.5.0.

Parent Defect ID:	EFA-5732	Issue ID:	EFA-5732
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	When firmware download is in progress, fabric delete command is accepted without an error.		
Condition:	If fabric delete command is submitted when firmware download is in progress, it fails.		
Workaround:	Allow firmware download process to complete. Status of the same can be checked using command efa inventory device firmware-download show --fabric {fabric name}		
Recovery:	Fabric can be deleted once the firmware download is completed		

Parent Defect ID:	EFA-6501	Issue ID:	EFA-6501
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Configuration Drift for VRF still shown in "cfg-in-sync" though its child configuration are drifted on SLX switch.		
Condition:	With below steps issue can be observed. 1) Create VRF/EPG having route target, static route and bgp configuration. 2) Introduce drift in VRF route target or static route or bgp configuration on SLX switch. 3) Update device from efa command "efa inventory device update --ip <device ip>" 4) Check device drift using efa command as "efa inventory drift-reconcile execute --ip <device ip>" 5) VRF shows as "cfg-in-sync" though its child configuration was drifted.		
Workaround:	None		
Recovery:	After drift and reconcile all EFA and device configuration will be in sync.		

Parent Defect ID:	EFA-7324	Issue ID:	EFA-7324
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Continuous create/delete of BGP peer-group and peer can finally cause CLI errors		
Condition:	When create and delete BGP peer/peer-group is repeatedly done in a loop. This will cause inventory does not have chance to update its DB with the current design so DB can be out of sync between inventory and tenant. When other events happen such as timer collection from inventory to sweep config to tenant, it can cause issues on tenant DB where CLI can fail.		

Parent Defect ID:	EFA-7324	Issue ID:	EFA-7324
Workaround:	Avoid such cycles of operations		
Recovery:	Delete the BGP peer/peer-group in problem and recreate them again.		

Parent Defect ID:	EFA-8090	Issue ID:	EFA-8090
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	When a fabric containing more than 15 newly registered devices is deployed using the CLI 'efa fabric configure', an attempt to add ports of any of these devices to a tenant within 5 minutes may fail. The error will indicate that the ports have not yet been registered in the fabric		
Condition:	Attempt to add device ports of a recently configured fabric to a tenant may fail with an error indication that the ports have not yet been registered in the fabric		
Workaround:	Wait for up to 5 minutes after deploying the fabric before adding ports to a tenant		
Recovery:	This is a transient error. Rerunning the port-add operation after a maximum wait time of 5 minutes will succeed		

Parent Defect ID:	EFA-8152	Issue ID:	EFA-8152
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	While graceful-restart(GR) updating with value TRUE and inflight transition triggered as a part of EFA rollover then update will continue as a part of inflight transition.		
Condition:	Update GR with value TRUE and perform EFA rollover on HA setup.		

Parent Defect ID:	EFA-8155	Issue ID:	EFA-8155
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	"cluster-client auto" is not configured under port channel for first reloaded device.		

Parent Defect ID:	EFA-8155	Issue ID:	EFA-8155
Condition:	Execute below steps to hit this condition 1) Create fabric on MCT paired device 2) Create Tenant/PO/VRF/EPG 3) Enable MM mode on both device 4) Perform EFA backup 5) Delete EPG/VRF/PO/Tenant 6) Delete fabric 7) Restore EFA backup 8) Reload device one by one After these steps check PO on both device, "cluster-client auto" will not configured on first reloaded device.		
Workaround:	Instead of reload device in Step (8), perform manual DRC using inventory CLI as "efa inventory drift-reconcile execute --ip <device ip> --reconcile" for each device.		

Parent Defect ID:	EFA-8257	Issue ID:	EFA-8257
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EFA is not able to detect drift for configuration like VRF/VE/VLAN/EVPN		
Condition:	Please follow below steps, 1) Create tenant/VRF/PO/EPG 2) As soon as EPG creation pushed configuration on device, remove them from device. 3) Check drift using inventory CLI as "efa inventory drift-reconcile execute --ip --device-ip <device ip>"		
Workaround:	As this is timing issue so we need to wait for 1 min before remove configurations from device.		
Recovery:	We need to delete EPG and recreate it again.		

Parent Defect ID:	EFA-8269	Issue ID:	EFA-8269
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG app-state moved to cfg-refresh-err after epg delete and admin up		

Parent Defect ID:	EFA-8269	Issue ID:	EFA-8269
Condition:	1) Configure clos fabric (Medium scale fabric) 2) Create tenant 3) Admin down the devices 4) Create port-channels, vrfs and epgs 5) Admin up the following devices Wait for the DRC to be success 6) Repeat step 3 Wait for the devices to put into maintenance mode 7) Create bgp peer-group and dynamic peers 8) Delete all epg's 9) Repeat step 5 10) Vrfs are getting deleted from admin up devices 11) EPG app-state move to cfg-refresh-err			
Recovery:	Delete the EPGs in cfg-refresh-err state and recreate them.		

Parent Defect ID:	EFA-8273	Issue ID:	EFA-8273
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG Update "vrf-add" operation gives success when EPG is in "vrf-delete-pending" state		
Condition:	Perform EPG Update "vrf-add" operation on an EPG in "vrf-delete-pending" state		
Workaround:	No workaround		
Recovery:	User needs to remove the VRF from EPG using EPG update "vrf-delete" operation before attempting the "vrf-add" operation.		

Parent Defect ID:	EFA-8315	Issue ID:	EFA-8315
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	User adds ports in empty EPG and immediately deletes them. The following adding ports into EPG can have error as duplicate entry		
Condition:	1) Add ports in empty EPG 2) Delete ports from epg right away 3) Add ports into EPG. Which can have error.		
Workaround:	After adding ports into EPG, wait certain time before trying to delete ports from EPG.		
Recovery:	Delete the EPG and recreate again.		

Parent Defect ID:	EFA-8322	Issue ID:	EFA-8322
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0

Parent Defect ID:	EFA-8322	Issue ID:	EFA-8322
Symptom:	EPG Update "anycast-ip-delete" operation gives different output/result when one of the EPG device is admin down		
Condition:	1) Create L3 EPG with anycast-ip/anycast-ipv6 2) Take one of EPG device administratively down 3) Bring device admin up which was taken down in previous step 4) While device is coming up administratively, try EPG Update "anycast-ip-delete" operation		
Workaround:	No workaround		
Recovery:	No recovery as such. Wait for device to be completely up before trying EPG Update "anycast-ip-delete" operation		

Parent Defect ID:	EFA-8334	Issue ID:	EFA-8334
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	system backup and restore causes epg state to be in cfg-refresh-err		
Condition:	Tenant DB and inventory DB needs time to be in sync. In a busy and scaled system this in sync can take much longer time to finish. Backup DB during DB un-sync window can cause system saves the DBs for tenant and inventory which are not synced yet and following restore will have issues.		
Workaround:	If there's a need to make system backup, please execute the backup after system have not made any new config for few minutes. It's needed for the inventory and tenant databases to be in sync before executing system backup. In a busy system the DB sync can take longer to finish.		
Recovery:	Delete the EPGs which report errors and recreate them.		

Parent Defect ID:	EFA-8335	Issue ID:	EFA-8335
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	system backup and restore causes following manual DRC has errors		
Condition:	Tenant DB and inventory DB needs time to be in sync. In a busy and scaled system this in sync can take much longer time to finish. Backup DB during this window will cause system saves the DB for tenant and inventory which is not synced yet and following restore will have issues.		

Parent Defect ID:	EFA-8335	Issue ID:	EFA-8335
Workaround:	If there's a need to make system backup, please execute the backup after system have not made any new config for few minutes. It's needed for the inventory and tenant databases to be in sync before executing system backup. In a busy system the DB sync can take longer to finish.		
Recovery:	Delete the epg or tenant with problem and recreate them.		

Parent Defect ID:	EFA-8443	Issue ID:	EFA-8443
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	For Tenant created with L3 port having multiple ip-address associated with it, "efa tenant show" will have repeated entries of that L3 port.		
Condition:	Steps to reproduce issue: 1) Assign multiple IPs to the physical port on SLX. 2) Create Tenant using same L3 port. 3) Check Tenant show output. L3 ports having multiple IPs will have repeated entry in the "efa tenant show" output.		
Workaround:	No workaround.		
Recovery:	Recovery can be done by removing all but one IP from the L3 port on SLX followed by an inventory device update.		

Parent Defect ID:	EFA-8465	Issue ID:	EFA-8465
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	The "efa inventory device firmware-download prepare add" command fails with "Please specify 'fullinstall' option in firmware download cmdline as GLIBC versions change".		
Condition:	Upgrading the SLX firmware from 20.1.2x to 20.2.x requires a 'fullinstall' firmware download in order to proceed.		
Workaround:	There is no workaround from EFA. The firmware download fullinstall must be carried out individually on each SLX device.		

Parent Defect ID:	EFA-8507	Issue ID:	EFA-8507
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	Certain vlans are missed in configuration when stacks are created in quick succession within a script with no delay.		

Parent Defect ID:	EFA-8507	Issue ID:	EFA-8507
Condition:	10 stack creations without much delay leads to missing configuration. Trunk Subport update is not generated from neutron. Issue is seen with only one controller and not seen when more delay is introduced between stack creations. Trunk also remains in DOWN state.		
Workaround:	Workaround is to have delay between stack creation.		
Recovery:	Remove the Trunk Parent port added to the VM and add it back again. e.g. Max-L2-ss3VirtIoVM2_Test1==> VM Name Max-L2-ss3VirtIoTrunkPort2_Test1 ==> Parent Port of the Sub Port that is down openstack server remove port Max-L2-ss3VirtIoVM2_Test1 Max-L2-ss3VirtIoTrunkPort2_Test1 openstack server add port Max-L2-ss3VirtIoVM2_Test1 Max-L2-ss3VirtIoTrunkPort2_Test1		

Parent Defect ID:	EFA-8512	Issue ID:	EFA-8512
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	On SLX there can have partial config of neighbor under router bgp. The "show running command router bgp" from SLX shows invalid command "neighbor pg1" (Assume the bgp-group name is pg1). There's no corresponding command to delete this.		
Condition:	It's found if issue netconf RPC to SLX device with BGP peer group delete operation which the peer-group does not exist, SLX will create the invalid "neighbor pg1".		
Workaround:	Under some admin-down device scenario, avoid delete the same bgp-peer more than once.		
Recovery:	On SLX use the following commands to get rid of the partial bgp-peer. SLX(config)# router bgp SLX(config-bgp-router)# neighbor pg1 peer-group SLX(config-bgp-router)# no neighbor pg1 peer-group		

Parent Defect ID:	EFA-8526	Issue ID:	EFA-8526
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	VRF Update "centralized-router-add" fail with error "[x, y] are MCT pair. Update the VRF with both devices together as centralized routers"		
Condition:	1) In a CLOS fabric setup with MCT pair of border-leafs, create VRF with routing-type as centralized and select MCT pair of border-leafs as centralized routers. 2) Remove one of the MCT pair border-leaf from the fabric 3) Add same/different border-leaf to the fabric and run fabric configure command 4) Wait for sometime and run VRF Update "centralized-router-add" operation to add newly added border-leaf as centralized router		

Parent Defect ID:	EFA-8526	Issue ID:	EFA-8526
Workaround:	Run VRF Update "centralized-router-add" operation and specify both nodes of MCT pair border-leafs as centralized routers.		
Recovery:	No recovery is required.		

Parent Defect ID:	EFA-8573	Issue ID:	EFA-8573
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	In few cases, networks in EPG will remain in cfg-in-sync state even if they are created with partial success topology (MCT pair with one admin-up device and one admin-down device).		
Condition:	The issue is seen with the below steps 1) Configure a fabric 2) Create Tenant 3) Create multi-homed portchannel 4) Bring one of the devices of the MCT pair(having the PO created in step 3) admin-down to create a partial success topology 5) Create EPGs on the partial success topology		
Recovery:	Bring all the devices in admin-up state. It should push all the configs on devices and everything will be in cfg-in-sync.		

Parent Defect ID:	EFA-8628	Issue ID:	EFA-8628
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	Tenant does not contain any Ports, state of EPG is "vrf-delete-pending" and EPG Update "vrf-delete" fail with error- "EPG cannot be updated with tenant having no ports"		
Condition:	Step1) Create Fabric and Tenant Step2) Create VRF using Routing type as Centralized and Border-leaf devices as Centralized Routers Step3) Create EPG using VRF created in Step2 with Physical Ports/Portchannel from leaf devices Step4) Delete all the devices from Fabric and/or Inventory Step5) Perform EPG Update "vrf-delete" operation		
Workaround:	No workaround is required.		
Recovery:	Delete and re-create EPG(s).		

Parent Defect ID:	EFA-8665	Issue ID:	EFA-8665
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1

Parent Defect ID:	EFA-8665	Issue ID:	EFA-8665
Symptom:	VRF configuration present on Border Leaf devices		
Condition:	1) Create Fabric and Tenant 2) Create VRF using Routing-type as "Centralized" and Border-leaf devices as Centralized-Routers 3) Create EPG(s) using VRF created in Step2 and Physical Ports/Portchannel from Leaf devices 4) Remove Leaf devices (having the Physical Ports/Portchannel used in the EPG) from Fabric/Inventory		
Workaround:	No workaround.		
Recovery:	Manually clean up VRF configuration from device		

Parent Defect ID:	EFA-8669	Issue ID:	EFA-8669
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EFA is not reachable after secondary node down/up of 8720. (kube-system pods are not responding)		
Condition:	This issue occurs only when the management cables are unplugged for under 10 seconds and plugged back in. What happens is that the EFA failover keepalived scripts, which does labeling and then switches the pod over to the other node, doesn't finish to completion causing the node labeling to not reflect right status where the VIP is assigned.		
Workaround:	Either keep the management ports unplugged for more than 20 seconds or reboot the tpvm.		
Recovery:	Reboot the tpvm.		

Parent Defect ID:	EFA-8701	Issue ID:	EFA-8701
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	EFA becomes unresponsive and the OS reports that there is no disk space.		
Condition:	Galera is technology used by mysql to keep data in sync between nodes. Occasionally, such a transfer between nodes can fail, resulting in a log file of the failed transaction, and some binary data. These accumulate in /apps/efadata/mysql. EFA will tar.gz these logs hourly and delete; but if there is a rapid accumulation, this process can be slow, and under extreme circumstances, not complete before the next hour starts. Eventually, the system is unable to proceed.		

Parent Defect ID:	EFA-8701	Issue ID:	EFA-8701
Workaround:	Delete files matching GRA_*.log in /apps/efadata/mysql. These are not needed by EFA.		
Recovery:	Delete files matching GRA_*.log in /apps/efadata/mysql. These are not needed by EFA. Kill all tar czvf processes working on GRA*.log files.		

Parent Defect ID:	EFA-8773	Issue ID:	EFA-8773
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	When the address family config is manually removed from the SLX followed by DRC, the DRC fails with the error - "Error: VRF Address Family not configured".		
Condition:	Steps to reproduce: 1. Configure a L3 EPG via EFA which results in VRF configuration on the SLX. 2. Remove the address-family configuration from SLX (this would remove bgp address-family configuration also). 3. Initiate the DRC to push back the configuration from EFA. 4. DRC fails with error "Error: VRF Address Family not configured".		
Workaround:	No workaround.		
Recovery:	Configure the address family on the switch manually and then trigger DRC. The configuration will be reconciled.		

Parent Defect ID:	EFA-8802	Issue ID:	EFA-8802
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Same epg will be allowed to update with different port properties value, however the first configured value will take effect		
Condition:	1. Create L3 epg e1 with port-group P1 port-group property switch port mode trunk 2. Update epg e1 by port-group-add P2 and port-group property switch port mode as trunk 3. Update same epg e1 by port-group-add P2 and different port-group property switch port mode trunk-no-default-native Above steps are applicable for all port properties like switch port mode , switchport-native-vlan-tagging, switchport-native-vlan, single-homed-bfd-session-type. No error is seen while updating with different port property		
Workaround:	Updating the epg with similar port property values since idempotency doesn't work		

Parent Defect ID:	EFA-8827	Issue ID:	EFA-8827
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	CEP port will remain in "cfg-refreshed" state even after the DRC is successful.		

Parent Defect ID:	EFA-8827	Issue ID:	EFA-8827
Condition:	1. Create EPG with CEP ports. 2. Configure "MM enable on reboot" on the SLX and modify any of the CEP port properties. 3. Reboot the SLX. 4. MM triggered DRC will reconcile all the port properties but the CEP will remain in cfg-refreshed state.		
Recovery:	CEP in cfg-refreshed state can be removed from the EPG and readded to the EPG.		

Parent Defect ID:	EFA-8848	Issue ID:	EFA-8848
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	When a configuration drift is introduced on a physical port on the SLX device, followed by a DRC on the device, the reconcile status displays "Portchannel" and "Vlan" reconciliation status as success even though the same were not drifted.		
Condition:	Below are the steps to reproduce the issue: 1) Create Fabric/Tenant using MCT paired SLX devices. 2) Create an EPG using the physical port (P1) from an MCT SLX device (D1). 3) Introduce drift on port (P1) on the SLX device(D1). 4) Perform DRC of the SLX device (D1). 5) DRC executed successfully and all configuration pushed on the SLX device. 6) DRC output displays "Portchannel" and "Vlan" reconcilation status as success even though the same were not drifted.		

Parent Defect ID:	EFA-8966	Issue ID:	EFA-8966
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	BGP peer group deletion fails when the deletion is attempted using a tenant not owning the BGP peer group.		
Condition:	Deletion of the BGP peer group which is not owned by an existing tenant.		
Workaround:	No workaround		

Parent Defect ID:	EFA-9009	Issue ID:	EFA-9009
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Creation or update of an EPG involving a port group comprising more than 50 ports and 100 or more VLANs may take more than 6 minutes and may fail		
Condition:	When a user sends a Tenant EPG command to create or update an EPG involving a port group comprising more than 50 ports and 100 or more VLANs, the command may fail with the error "Error : EPG: <epg-name> Save for devices failed"		

Parent Defect ID:	EFA-9009	Issue ID:	EFA-9009
Workaround:	Split such EPG create or update request to add not more than 50 VLANs at a time.		
Recovery:	None.		

Parent Defect ID:	EFA-9400	Issue ID:	EFA-9400
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.4
Symptom:	When the EPG is in "port-group-delete-pending" state, the subsequent "port-group-delete" operations will not clean up any configurations from the admin up devices.		
Condition:	The issue can be reproduced with the below steps: 1. Create an EPG with CEEP port channels pg1, pg2. 2. Admin down one of the MCT nodes. 3. Deletion of pg1 from the EPG (created in step 1) will clean up all configs from the admin up device and the EPG moves to the "port-group-delete-pending" state. 4. Deletion of pg2 from the EPG will not clean any of the configurations from admin up device since the EPG is in the "port-group-delete-pending" state.		
Workaround:	Perform a single port-group-delete operation with all the port-groups belonging to the EPG, instead of multiple port-group-delete operations with each iteration containing some of the port-groups belonging to the EPG		
Recovery:	Bring all the devices in admin-up state and then perform the port-group-delete operation on the EPG which will clean all configs on both the nodes		

Parent Defect ID:	EFA-9443	Issue ID:	EFA-9443
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	resilient-hash max-path is seen on the vrf even though it was deleted from the VRF (when the device was down) using VRF update rh-max-path-delete operation followed by the admin up of the device.		
Condition:	Below are the steps to reproduce the issue: 1. Configure Tenant and VRF with rh-max-path. 2. Create EPG using VRF created in step1. 3. Admin-down one of the MCT devices. 4. Execute VRF Update rh-max-path-delete operation. 5. Bring the admin down device administratively up.		

Parent Defect ID:	EFA-9443	Issue ID:	EFA-9443
Workaround:	1. Execute VRF Update rh-max-path-add to add resilient-hash max-path again to the VRF. 2. Execute VRF Update rh-max-path-delete to remove resilient-hash max-path from the VRF.		
Recovery:	No recovery is required		

Parent Defect ID:	EFA-9451	Issue ID:	EFA-9451
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.3
Symptom:	In a brownfield deployment with the existing/stale MCT cluster-client configuration imported into EFA, the EPG creation fails when the MCT cluster-client ID conflicts with the MCT cluster-client ID already consumed by an out-of-band created MCT cluster-client.		
Condition:	Below are the probable steps to reproduce the issue: 1. Dual-homed (MCT) SLX has a brownfield MCT cluster-client configuration which is imported into EFA DB during "efa fabric configure" on the SLX. 2. Create tenant on the fabric configured in step1. 3. Create multi-homed port channels (which are already configured as MCT cluster-clients on the SLX) under the ownership of the tenant created in step 2. 4. Create EPG using the multi-homed port-channels and the same fails with the error "Port :<slx-mgmt-ip>:port-channel:<po-id> ClientID : <1000+po-id> already configured, conflicting with tenant-service generated ID <po-id>".		
Workaround:	No workaround		
Recovery:	Below are the recovery steps: 1. Delete the "cluster" configuration from one of the MCT SLX nodes and perform "inventory device update" for the SLX. 2. Delete the "cluster" configuration from the other MCT SLX node and perform "inventory device update" for the SLX. 3. Execute DRC for both the MCT SLX nodes.		

Parent Defect ID:	EFA-9467	Issue ID:	EFA-9467
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When a BGP peer-group create with an invalid MD5 password is attempted on an SLX device via EFA, the creation will fail with a valid error from SLX, resulting in the rollback of the failed operation which further results in a stale/partial configuration "neighbor <peer-group-name>" on the SLX.		

Parent Defect ID:	EFA-9467	Issue ID:	EFA-9467
Condition:	Below are the steps to reproduce the issue: 1. Configure a CLOS/Non-CLOS fabric. 2. Create a tenant on the fabric configured in step 1. 3. Create a BGP peer-group (under the ownership of the tenant created in step 2) with an invalid md5 password, which fails with a valid error from SLX device. 4. Failure of (3) results in the "rollback" operation resulting in the removal of peer-group configuration from the SLX device, which further results in a stale/partial SLX configuration "neighbor <peer-group-name>".		
Workaround:	No workaround.		
Recovery:	On the SLX, use the following commands to remove the partial bgp peer-group configuration. SLX(config)# router bgp SLX(config-bgp-router)# neighbor pg1 peer-group SLX(config-bgp-router)# no neighbor pg1 peer-group		

Parent Defect ID:	EFA-9487	Issue ID:	EFA-9487
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When bgp static peers are created (on a partial success topology) with the target device being admin down device followed by admin up of the admin down device, the dev-state/app-state continue to be not-provisioned/cfg-refreshed instead of provisioned/cfg-in-sync.		
Condition:	This issue can be reproduced with the below steps. 1. Admin down one of the SLX devices of an MCT pair. 2. Create BGP static peers with md5-password with the admin down SLX device as the target device. 3. Admin up the SLX device (which was admin down as mentioned in step 1). 4. After admin up, the BGP static peers' dev-state/app-state moves to not-provisioned/cfg-refreshed instead of it being provisioned/cfg-in-sync.		
Workaround:	Create the bgp static peers with md5-password when all the devices are in admin-up state.		
Recovery:	1. Bring all the devices to the admin up state. 2. Delete the BGP static peers which were created with md5-password when one of the devices was admin down. 3. Recreate the BGP static peers which were created with md5-password when one of the devices was admin down.		

Defects Closed without Code Changes

The following defects were closed in Extreme Fabric Automation 2.5.0.

Parent Defect ID:	EFA-5841	Issue ID:	EFA-5841
Reason Code:	Working as Designed	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	When firmware download is in progress, tenant create command is accepted without an error.		
Condition:	If tenant commands are submitted when firmware download is in progress, it results in erroneous configuration and some configurations may miss.		
Workaround:	Allow firmware download process to complete. Status of the same can be checked using command efa inventory device firmware-download show --fabric {fabric name}		
Recovery:	Tenant commands can be submitted after the firmware download is completed		

Parent Defect ID:	EFA-5874	Issue ID:	EFA-5874
Reason Code:	Already Implemented	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	On device registration, the IP of the EFA system is recorded in the logging entry on the device so logs can be forwarded to the EFA system for notification. When the EFA system is backed up and restored on another system with a different IP, the old IP of the EFA system is still present on the devices and the devices will continue to forward logs to the old EFA IP.		
Workaround:	Users will have to manually login to each devices and remove the logging entry for the old EFA IP.		

Parent Defect ID:	EFA-7592	Issue ID:	EFA-7592
Reason Code:	Working as Designed	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.2
Symptom:	"dev-state/app-state" moved to not-provisoned/cfg-ready		
Condition:	1) Configure non-clos fabric 2) Create tenant, vrf, epg 3) Admin down device 4) Create multiple epg's , delete an existing epg 5) Manually delete vrf from admin down device 6) Admin up device 7) After admin up, for epg which is in delete-pending theapp-state moved to cfg-ready		

Parent Defect ID:	EFA-7592	Issue ID:	EFA-7592
Workaround:	wait for few minutes after epg delete, before admin up of the device.		
Recovery:	force delete the EPGs in question and recreate them.		

Parent Defect ID:	EFA-8319	Issue ID:	EFA-8319
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	PO creation failed with error "Devices are not MCT Pairs".		
Condition:	Please follow below steps, 1) Create fabric/tenant/PO/EPG 2) Take EFA backup 3) Delete EPG/PO/tenant/fabric 4) Restore EFA backup taken in step (2) 5) Delete tenant from which was created before backup 6) Create same tenant again 7) Create PO under same tenant		
Workaround:	As after restore MCT peer details are Nil so we need to perform DRC after restore taken backup. After step (4) above, we need to perform DRC using inventory CLI as efa inventory drift-reconcile execute --ip <device ip 1> --reconcile efa inventory drift-reconcile execute --ip <device ip 2> --reconcile		

Parent Defect ID:	EFA-8754	Issue ID:	EFA-8754
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Shutting both mgmt and RME ports can occasionally cause EFA's kubernetes deployment to become stuck in an initializing state.		
Condition:	Turning off and on both mgmt and RME ports.		
Recovery:	Restart kubernetes on both nodes of the cluster. As root user: $ systemctl restart k3s		

Parent Defect ID:	EFA-8967	Issue ID:	EFA-8967
Reason Code:	Working as Designed	Severity:	S4 - Low
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.2
Symptom:	Error during epg update ctag-range-add operation if different anycastIP is provided for new ctag range with same l2-vni already configured in EPG. Ex. Error : EPG: e1(11:11.11.11.1/24) shares network name with EPG: e1(10:10.10.10.1/24) should have same Anycast IP		

Parent Defect ID:	EFA-8967	Issue ID:	EFA-8967
Condition:	1. Create L3 EPG epg1 with ctag-1, l2vni-1, anycastIP1 2. Update EPG epg1 with ctag-2, l2vni-1, anycastIP2 For the same l2vni, anycastIP must be the same, the same condition will be verified as part of epg validation which provided above-mentioned error.		
Workaround:	1. For same l2vni, anycastIP must be same. 2. Different l2-vni can be used for using different anycastIP.		
Recovery:	NA		

Parent Defect ID:	EFA-9045	Issue ID:	EFA-9045
Reason Code:	Insufficient Information	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.2
Symptom:	1. "app-state" of some of the VRFs is shown as "cfg-refreshed" in "efa tenant vrf show" output. 2. The same VRFs whose "app-state" is shown as "cfg-refreshed" are shown as "unstable" VRFs in the "efa tenant epg show" output.		
Condition:	1. Create an 18 node CLOS fabric. 2. Create multiple tenants (e.g. 14) tenants under the fabric created in step 1. 3. Create multiple VRFs (e.g. 400+) under the tenant created in step 2. 4. Create multiple EPGs using the VRFs created in step 3. 5. Check the "efa tenant vrf show" output to know the "app-state" of the VRFs.		
Workaround:	No workaround.		
Recovery:	1. Check the configuration drift per SLX device using the CLI "efa inventory drift-reconcile execute --ip <slx-device-ip>" to identify if the VRFs (which are in the cfg-refreshed state) are shown in the drift output. 2. If the VRFs are shown in the drift output for a given SLX device, then delete the VRFs from that particular SLX device. 3. Perform DRC (Drift and Reconcile) for the SLX device from which the unstable VRFs were deleted. 4. Steps 3 and 4 need to be done for all the SLX devices.		

Parent Defect ID:	EFA-9227	Issue ID:	EFA-9227
Reason Code:	Working as Designed	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	After the service restart, the inflight operations i.e "VRF update max-path-delete" or "VRF update rh-max-path-delete" are rolled forward instead of rollback.		
Condition:	When "VRF Update max-path-delete" or "VRF Update rh-max-path-delete" operation is executed and the tenant service gets restarted while the operation is in progress.		

Parent Defect ID:	EFA-9227	Issue ID:	EFA-9227
Workaround:	No workaround		
Recovery:	No recovery		

Parent Defect ID:	EFA-10075	Issue ID:	EFA-10075
Reason Code:	Already Reported	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.6
Symptom:	When the EFA created port-channel is deleted (manually on the SLX), re-created (manually on the SLX) without member ports, with some port-channel config drift (e.g. po speed), then the DRC doesn't identify any config drift.		
Condition:	Below are the steps to reproduce the issue: 1. Create Port-channel on the SLX using EFA 2. Delete the port-channel from the SLX manually 3. Re-create port-channel on the SLX manually with drifted speed configuration 4. Perform DRC without reconciliation		
Recovery:	Delete the port-channel manually on SLX and trigger DRC from EFA		

Open Defects

The following defects are open in Extreme Fabric Automation 2.5.0.

Parent Defect ID:	EFA-10126	Issue ID:	EFA-10126
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Users might see a failure message when doing a node replacement when running the installer in GUI mode, when the progress is around 18% although the node replacement proceeds		
Condition:	This error message is not expected and is not consistently seen, however this issue is not a functional issue but more of a display issue		
Workaround:	Since the installation and node replacement proceeds fine and succeeds, users can just ignore the briefly seen failed message on the installer. This will be fixed in an upcoming release.		

Parent Defect ID:	EFA-10121	Issue ID:	EFA-10121
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	TPVM and Apps/EFA running in TPVM, will be removed after copy default-config startup-config" and "reload system" commands on SLX.		

Parent Defect ID:	EFA-10121	Issue ID:	EFA-10121
Condition:	If copy default-config startup-config is run on the SLX device that is hosting a TPVM, the TPVM and applications/EFA running within, will be removed after "reload system"		

Parent Defect ID:	EFA-10121	Issue ID:	EFA-10121
Workaround:	Avoid using "copy default-config startup-config"		
Recovery:	This recovery procedure assumes multi-node EFA deployment on SLX1/TPVM1 and SLX2/TPVM2.. TPVM1 is the TPVM that was removed on SLX1, and SLX2/TPVM2 is the second device and now running active EFA. On SLX1 perform the following steps. 1. Reinstall TPVM1 using config mode with the same configuration as before. 2. On SLX2/TPVM2, setup trusted peer with TPVM1. "SLX2(config-tpvm-TPVM)# trusted-peer ip TPVM1 password <password>" 3. On TPVM2, perform efa upgrade with node replacement using the following steps. SLX2# efa deploy --graphics no Starting "efa deploy"... Step 1: Checking if TPVM is deployed... Step 2: Get IP Address assigned to TPVM to deploy EFA IP Address of the TPVM TPVM2 Step 3: Checking for EFA packages in /efaboot directory Step 4: Deploying EFA package efa-2.5.0.tar.gz on TPVM2 Step 5: Checking for EFA Stack... Previous Stack found Are you sure you want to re-deploy EFA? (yes/no) yes Step 6: Re-deploying EFA Would you like to configure additional IP address(es) for the HA health ping check? (yes/no) no Would you like to configure additional management IP networks? (yes/no) no Would you like to configure additional management IP network routes? (yes/no) no Please choose: 1 Multi Node Build Upgrade 2 Multi Node Build Upgrade With Node Replacement 2 Enter the replacement peer node (IP/Hostname): TPVM1 TPVM1-IP server is reachable... You have chosen: - to redeploy EFA at version 2.5.0 build GA - with peer TPVM1 and VIP EFA-VIP - with node replacement Do you wish to proceed? (yes/no) yes ... ... Waiting for EFA services Waiting for EFA containers to start Extreme Fabric Automation Stack has been upgraded successfully. SLX2#		

Parent Defect ID:	EFA-10121	Issue ID:	EFA-10121
	Once efa upgrade with node replacement is complete, verify on both nodes, efa is running and services have started using efactl status.		

Parent Defect ID:	EFA-10115	Issue ID:	EFA-10115
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Redeployment of EFA 2.5.0 fails on a multi-node when one of the nodes in the cluster is changed.		
Condition:	After un-deployment of multi-node EFA, if a fresh deployment is triggered after changing one of the nodes or its IP in the cluster, then installation fails.		
Workaround:	Edit /etc/fstab file and remove any mount point entries of old IPs from both the nodes in the multi-node setup.		

Parent Defect ID:	EFA-10110	Issue ID:	EFA-10110
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	EFA fabric and tenant operations are not blocked when (manual) DRC operation is triggered and in-progress.		
Condition:	DRC operations may fail/timeout with fabric/tenant operations taking longer time to complete.		
Workaround:	Do not run fabric configure or tenant operations when manual DRC for a device is in progress.		

Parent Defect ID:	EFA-10109	Issue ID:	EFA-10109
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Global BFD timer values will be applied to multi-hop BFD sessions in multi-rack setups.		
Condition:	This behavior is seen only for multi-rack, multi-hop BFD sessions.		
Workaround:	Set required bfd timer values under appropriate peer-group using SLX CLI		
Recovery:	Set required bfd timer values under appropriate peer-group using SLX CLI		

Parent Defect ID:	EFA-10099	Issue ID:	EFA-10099
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0

Parent Defect ID:	EFA-10099	Issue ID:	EFA-10099
Symptom:	When the md5-password is updated on an already provisioned fabric, the existing tenant vrf backup routing bgp neighbours will be updated with the new md5-password followed by the clearing of the bgp neighbours. Updation of md5-password and clearing of the corresponding bgp neighbors happens one SLX device at a time, hence resulting in the session being down time till the process is complete for both the devices of the MCT pair.		
Condition:	1. Configure fabric with the fabric setting backup routing enabled and with the md5-password fabric setting 2. Configure tenant under the fabric 3. Configure VRF and L3 EPG (using the VRF and under the ownership of the tenant), which results in the creation of the backup routing bgp neighbors (for the tenant vrf) using the md5-password provided at the fabric setting 4. Update md5-password on the already provisioned fabric followed by "fabric configure"		
Workaround:	No workaround		

Parent Defect ID:	EFA-10093	Issue ID:	EFA-10093
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Deletion of the VLAN/BD based L3 EPGs in epg-delete-pending state will result in creation and then deletion of the VLAN/BD on the admin up device where the VLAN/BD was already removed		
Condition:	Issue occurs with the below steps: 1. Create L3 EPG with VLAN/BD X on an MCT pair 2. Admin down one of the devices of the MCT pair 3. Delete the L3 EPG. This results in the L3 configuration removal (corresponding to the L3 EPG getting deleted) from the admin up device and no config changes happen on the admin down device and the EPG transits to epg-delete-pending state 4. Admin up the device which was made admin down in step 2 5. Delete the L3 EPG which transited to epg-delete-pending state in step 3		
Recovery:	Not needed		

Parent Defect ID:	EFA-10067	Issue ID:	EFA-10067
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0

Parent Defect ID:	EFA-10067	Issue ID:	EFA-10067
Symptom:	In a node replacement scenario, the standby node will not have a TPVM configured nor running. When the "efa inventory device tpvm-upgrade" command is run against this replacement node, the TPVM deployment and upgrade to the new TPVM version fails.		
Condition:	The TPVM is neither configured nor running on the switch.		

Parent Defect ID:	EFA-10064	Issue ID:	EFA-10064
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	During the required fullinstall firmware download of SLXOS from 20.2.3f to 20.3.2a the TPVM configuration through exec-mode commands are not converted to the running-config.		
Condition:	When a firmware download is run using the fullinstall option.		

Parent Defect ID:	EFA-10063	Issue ID:	EFA-10063
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Deleting device from EFA Inventory would not bring up interface to admin state 'up' after unconfiguring breakout configuration		
Condition:	This condition occurs when there is breakout configuration present on device that is being deleted from EFA Inventory		
Workaround:	Manually bring the admin-state up on the interface, if required		
Recovery:	Manually bring the admin-state up on the interface, if required		

Parent Defect ID:	EFA-10062	Issue ID:	EFA-10062
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Removing device from Inventory would not clean up breakout configuration on interfaces that are part of port channels.		
Condition:	This condition occurs when there is breakout configuration present on device that is being deleted from EFA Inventory, such that those breakout configurations are on interfaces that are part of port-channels		

Parent Defect ID:	EFA-10062	Issue ID:	EFA-10062
Workaround:	Manually remove the breakout configuration, if required.		
Recovery:	Manually remove the breakout configuration, if required.		

Parent Defect ID:	EFA-10048	Issue ID:	EFA-10048
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	EPG: epgev10 Save for devices failed When concurrent EFA tenant EPG create or update operation is requested where the commands involve large number of vlans and/or ports, one of them could fail with the error "EPG: <epg-name> Save for devices Failed".		
Condition:	The failure is reported when concurrent DB write operation are done by EFA Tenant service as part of the command execution.		
Workaround:	This is a transient error and there is no workaround. The failing command can be executed once again and it will succeed.		
Recovery:	The failing command can be rerun separately and it will succeed.		

Parent Defect ID:	EFA-10041	Issue ID:	EFA-10041
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	The "efa inventory device tpvm-upgrade execute" command to upgrade to a new tpvm version will result in a failed tpvm upgrade where the previous tpvm image will be rolled back and restored.		
Condition:	The "trusted peer" TPVM running-configuration has not been applied to either SLX device hosting the TPVMs installed with an EFA multi-node deployment.		
Workaround:	Ensure the "trusted peer" TPVM running-configuration is configured on one of the SLX devices.		
Recovery:	Configure the "trusted peer" TPVM running-configuration and rerun the TPVM upgrade command.		

Parent Defect ID:	EFA-10026	Issue ID:	EFA-10026
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	'efa inventory device interface unset-fec' command will set the fec mode to 'auto-negotiation' instead of removing fec configuration.		
Condition:	Once fec mode is set on the interface, the configuration cannot be removed. 'efa inventory device interface unset-fec' command will set the fec mode to 'auto-negotiation' instead of removing fec configuration. This is because 'no fec mode' command is no longer supported on SLX.		

Parent Defect ID:	EFA-10026	Issue ID:	EFA-10026
Workaround:	Default value for fec-mode is 'auto-negotiation' and will show up as-is in the output of 'show running-config'. Users can set a different value using 'efa inventory device interface set-fec', if required.		
Recovery:	Default value for fec-mode is 'auto-negotiation' and will show up as-is in the output of 'show running-config'. Users can set a different value using 'efa inventory device interface set-fec', if required.		

Parent Defect ID:	EFA-10018	Issue ID:	EFA-10018
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Deployment may fail in the step "Checking default gateway reachability on all nodes" due to network reachability issue with gateway		
Condition:	Network reachability issue with gateway will cause deployment failure		
Workaround:	Prior to performing the deployment, login to the new tpvm and ping the gateway until it is reachable. Once it is, the deployment will then run fine.		
Recovery:	Cancel the deployment & Login to the redeployed TPVM. Ping the gateway IP until it is connected. Once it is, the deployment will run fine.		

Parent Defect ID:	EFA-10016	Issue ID:	EFA-10016
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	BGP Peer creation fails with the error "failed to fetch BgpAddressFamily data" because of the intermittent connectivity loss of EFA with SLX. Rollback also failed leaving the stale config on SLX.		
Condition:	1. Create tenant, po, vrf and epg 2. Create bgp peer group 3. Create bgp peers Create fails because of an intermittent connection issue		
Workaround:			
Recovery:	1. Delete the peer using force option so all stale config is removed. 2. Recreate the bgp peer		

Parent Defect ID:	EFA-9990	Issue ID:	EFA-9990
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	EPG update ctag-range-add operation with the existing ctag-range (i.e. ctag1, ctag2) and modified native vlan (ctag2) succeeds without any effect		

Parent Defect ID:	EFA-9990	Issue ID:	EFA-9990
Condition:	Below are the steps to reproduce the issue: 1. Create Endpoint group with ctag1, ctag2 and native vlan as ctag1 2. Update the Endpoint group (created in step 1) using ctag-range-add operation with the same set of ctags (i.e. ctag1, ctag2) and different native VLAN ctag2		
Workaround:	If user intends to modify the native vlan from ctag1 to ctag2 in an EPG, then the user will need to remove ctag1 (using ctag-range-delete) from the EPG and add ctag2 (using ctag-range-add) as native vlan to the EPG		

Parent Defect ID:	EFA-9988	Issue ID:	EFA-9988
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	If the SLX device is in firmware-download-in-progress state, then the DRC (drift and reconcile) fails, But the failure reason is not shown in the DRC output		
Condition:	Trigger drift and reconcile from EFA for the SLX device when the SLX device is in the firmware-download-in-progress-state		
Workaround:	Wait for the firware-download to complete and then execute DRC		

Parent Defect ID:	EFA-9968	Issue ID:	EFA-9968
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	DRC fail with error "Delete and Update operations are not supported in a singled transaction. Please try them individually."		
Condition:	Below are the steps to reproduce: 1. Create Tenant 2. Create VRF with max-path and graceful-restart-enable 3. Create EPG using VRF created in step 1 4. Take one of the SLX devices to administratively down state 5. Perform VRF Update max-path-add operation to add a different max-path value 6. Perform VRF Update graceful-restart-update to disable graceful-restart 7. Admin up the SLX device which was made administratively down in step 4 and wait for DRC to complete		
Workaround:	No workaround		
Recovery:	Below are the steps to recover: 1. Perform VRF Update graceful-restart-update to enable graceful-restart 2. Admin up the SLX device which was made administratively down and wait for DRC to complete 3. Perform VRF Update graceful-restart-update to disable graceful-restart		

Parent Defect ID:	EFA-9966	Issue ID:	EFA-9966
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	In 2.4.x, adding the same ports/pos with different native vlan was allowed without any validation, this leads to two possible issues 1. Adding the same ports/PO to multiple epgs with different native-vlan VLAN results in the initial input native vlan being over written with the latest native vlan value onto the SLX 2. Adding the same port/po to multiple epgs, the first epg being created with native vlan and the second epg being created without native vlan, will result in removal of native vlan configured as part of the first epg The same inconsistency gets carried over from 2.4.x to 2.5.0		
Condition:	Case1: 1. Create fabric, tenant 2. Create PO1 with Device1Port1 3. Create PO2 with Device1Port2 4. Create EPG1 with PO1, PO2, CTAG1, CTAG2 with "switchport mode trunk", and native vlan as CTAG1 5. Create EPG2 with PO2, CTAG3, CTAG4 with native vlan CTAG3 6. On the SLX, PO2 initial native vlan CTAG1 gets replaced with CTAG3, while PO1 will have the initial native vlan CTAG1 Case2: 1. Create fabric, tenant. 2. Create PO1 with Device1Port1 3. Create PO2 with Device1Port2 4. Create EPG1 with PO1, PO2, CTAG1, CTAG2 with "switchport mode trunk" and native vlan as CTAG1 5. Create EPG2 with PO2 CTAG3, CTAG4 without native vlan 6. On the SLX, PO2 initial native vlan CTAG1 gets replaced with the default value, while PO1 will have the initial native vlan CTAG1		
Workaround:	The workaround is not available.		
Recovery:	1. Identify EPGs with common PO/port with conflicting native vlan 2. Delete conflicting ctag (that needs to be configured as native vlan but isnt configured as one) using EPG update "ctag-range-delete" operation 3. Add ctag back with correct native vlan (to be configured on SLX) using EPG update ctag-range-add" operation		

Parent Defect ID:	EFA-9952	Issue ID:	EFA-9952
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	network-property delete failed When concurrent EFA tenant EPG delete operations are requested where the commands involve large number of vlans and/or ports, one of them could fail with the error "EPG network-property delete failed"		
Condition:	The failure is reported when concurrent DB write operation are done by EFA Tenant service as part of the command execution.		

Parent Defect ID:	EFA-9952	Issue ID:	EFA-9952
Workaround:	This is a transient error and there is no workaround. The failing command can be executed once again and it will succeed.		
Recovery:	The failing command can be rerun separately and it will succeed		

Parent Defect ID:	EFA-9945	Issue ID:	EFA-9945
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Make config changes of channel group mode on SLX for the interface belongs to some PO. Issue DRC from EFA. Drift and reconcile will not detect this change and correct it.		
Condition:	Steps to reproduce: Create tenant and PO Modify channel group mode on SLX config t interface Ethernet 0/17-18 no channel-group channel-group 111 mode passive type standard Update inventory Manual DRC		
Workaround:	Avoid making manual changes on device for channel group mode. This will be supported in later release.		

Parent Defect ID:	EFA-9944	Issue ID:	EFA-9944
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	If port channel speed is modified on the device. In the mean time member ports are deleted from PO. Issue DRC from EFA, the DRC will fail with error and PO speed is not get reconciled: "10.20.246.3:ethernet:0/19 to Portchannel po2 failed due to netconf rpc [error] %Error: Port-channel should be admin down for speed to be configured"		
Condition:	Steps to reproduce: 1. Create tenant and PO 2. Modify po speed on SLX config t no int po 112 interface Port-channel 112 speed 100 no shutdown 3. Update inventory 4. Manual DRC		

Parent Defect ID:	EFA-9944	Issue ID:	EFA-9944
Workaround:	On device avoid making PO speed change and remove member ports from PO at the same time. EFA will have support in later release.		
Recovery:	Delete the PO from device and issue DRC.		

Parent Defect ID:	EFA-9941	Issue ID:	EFA-9941
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.3
Symptom:	EPG create with a CEP port (already used in another EPG) fails with the NETCONF RPC Error "NOTAKNOWNRESOURCEID"		
Condition:	1. Create fabric and tenant. 2. Create EPG1 using CEP port. 3. Create EPG2 using the same CEP port.		
Recovery:	Delete the EPG using the common CEP port and recreate the EPG		

Parent Defect ID:	EFA-9932	Issue ID:	EFA-9932
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	System API throws a 502 error after a restore is run.		
Condition:	After executing 'efa system restore', triggering any system API sometimes yield a 502 response.		
Workaround:	Run 'efa status' to confirm EFA is running and re-execute the commands.		

Parent Defect ID:	EFA-9930	Issue ID:	EFA-9930
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Periodic backup happens according to the system timezone.		
Condition:	If the nodes in HA are not configured in the same timezone, then periodic backup is scheduled according to the timezone of the active node. When a failover happens, the schedule is changed to the timezone of the new active node.		
Workaround:	Configure the same timezone on both the nodes in a multi-node installation		

Parent Defect ID:	EFA-9906	Issue ID:	EFA-9906
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0

Parent Defect ID:	EFA-9906	Issue ID:	EFA-9906
Symptom:	When concurrent EFA tenant EPG create or update operation is requested where the commands involve large number of vlans and/or ports, one of them could fail with the error "EPG: <epg-name> Save for Vlan Records save Failed".		
Condition:	The failure is reported when concurrent DB write operation are done by EFA Tenant service as part of the command execution.		
Workaround:	This is a transient error and there is no workaround. The failing command can be executed once again and it will succeed.		
Recovery:	The failing command can be rerun separately and it will succeed.		

Parent Defect ID:	EFA-9874	Issue ID:	EFA-9874
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When EPG is in the "anycast-ip-delete-pending" state and the user performs "epg configure", it will succeed without actually removing anycast-ip from SLX.		
Condition:	Below are the steps to reproduce the issue: 1) Configure EPG with VRF, VLAN and anycast-ip (ipv4/ipv6) on a single rack Non-CLOS fabric. 2) Bring one of the devices to admin-down. 3) EPG Update anycast-ip-delete for anycast-ip ipv4 or ipv6. This will put EPG in "anycast-ip-delete-pending" state. 4) Bring the admin-down device to admin-up. 5) In this state, the only allowed operations on EPG are "epg configure" and EPG update "anycast-ip-delete". 6) Perform "epg configure --name <epg-name> --tenant <tenant-name>".		
Workaround:	No workaround.		
Recovery:	Perform the same anycast-ip-delete operation when both devices are admin-up.		

Parent Defect ID:	EFA-9813	Issue ID:	EFA-9813
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.3
Symptom:	When doing RMA of device the port connections for the new device must be identical.		
Condition:	New device's port connections were not identical to the device being RMAed.		
Workaround:	When doing RMA of device the port connections for the new device must be identical.		

Parent Defect ID:	EFA-9799	Issue ID:	EFA-9799
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0

Parent Defect ID:	EFA-9799	Issue ID:	EFA-9799
Symptom:	'efa status' response shows standby node status as 'UP' when node is still booting up		
Condition:	If SLX device is reloaded where EFA standby node resides, then 'efa status' command will still show the status of standby as UP.		
Workaround:	Retry the same command after sometime.		

Parent Defect ID:	EFA-9758	Issue ID:	EFA-9758
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When user modifies the remote-asn of BGP peer out of band, drift and reconcile is not reconciling the intended remote-asn of BGP peer configuration.		
Condition:	Issue will seen if the user modifies the remote ASN of BGP peer through out of band means, DRC is not reconciling the remote ASN.		
Workaround:	Login to the device where the remote ASN is modified and revert it back to what EFA has configured.		
Recovery:	Revert the remote ASN of BGP peer on the device through SLX CLI to what EFA has configured previously.		

Parent Defect ID:	EFA-9674	Issue ID:	EFA-9674
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.2
Symptom:	Creation and deletion of stacks can result in failure. Network create fails as the previous network with same VLAN is not deleted.		
Condition:	Network is deleted and created in quick succession. Since the EFA processing takes time to delete the network at EFA, another call made for network create with same vlan id is also processed. This network create call will end in failure.		
Workaround:	Add delay between delete and create of stacks to allow more time for efa processing.		
Recovery:	Cleanup and recreate the failed network/stack at openstack		

Parent Defect ID:	EFA-9669	Issue ID:	EFA-9669
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Network and router creation during EFA HA fail-over, stale entries Steps: Started script/stack to create networks and routers Did HA fail over when network/router creation was in progress Few network and router creations failed.		

Parent Defect ID:	EFA-9669	Issue ID:	EFA-9669
Condition:	EFA HA Failover during stack creation can result in failed network and router entries at OpenStack. EFA services will be down for few minutes during HA failover, resulting in the failures.		
Workaround:	No workaround. This is expected behavior during EFA HA failover.		
Recovery:	Delete and recreate the network/router entries at OpenStack after EFA HA failover is complete. Use 'efa-health show' to check EFA HA status.		

Parent Defect ID:	EFA-9645	Issue ID:	EFA-9645
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When the fabric setting is updated with this particular password "password$\n", md5 password doesn't get configured on the backup routing neighbors that was already created.		
Condition:	1. Configure fabric 2. Create tenant, po, vrf and epg 3. Update fabric setting with "password$\n" and configure fabric 4. MD5 password is not configured on backup routing neighbors under BGP address family ipv4/ipv6 vrf		
Workaround:	Update the fabric setting with any other password combination that does not include "$\n" combination.		
Recovery:	Update the fabric setting with any other password combination that does not include "$\n" combination.		

Parent Defect ID:	EFA-9591	Issue ID:	EFA-9591
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When certain BGP sessions are not in ESTABLISHED state after clearing the BGP sessions as part of fabric configure, we see this issue.		
Condition:	This condition was seen when "efa fabric configure --name <fabric name>" was issued after modifying the MD5 password.		
Workaround:	Wait for BGP sessions to be ready. Check the status of BGP sessions using "efa fabric topology show underlay --name <fabric name>"		
Recovery:	Wait for BGP sessions to be ready. Check the status of BGP sessions using "efa fabric topology show underlay --name <fabric name>"		

Parent Defect ID:	EFA-9576	Issue ID:	EFA-9576
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0

Parent Defect ID:	EFA-9576	Issue ID:	EFA-9576
Symptom:	Deletion of the tenant by force followed by the recreation of the tenant and POs can result in the error "Po number <id> not available on the devices".		
Condition:	Below are the steps to reproduce the issue: 1. Create tenant and PO. 2. Delete the tenant using the "force" option. 3. Recreate the tenant and recreate the PO in the short time window.		
Workaround:	Avoid performing tenant/PO create followed by tenant delete followed by the tenant and PO recreate in the short time window.		
Recovery:	Execute inventory device prior to the PO creation.		

Parent Defect ID:	EFA-9570	Issue ID:	EFA-9570
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Add Device Failed because ASN used in border leaf showing conflict		
Condition:	If there are more than one pair of Leaf/border leaf devices then devices which are getting added first will get the first available ASN in ascending order and in subsequent addition of devices if one of device is trying to allocate the same ASN because of brownfield scenario then EFA will throw an error of conflicting ASN		
Workaround:	Add the devices to fabric in following sequence 1)First add brownfield devices which have preconfigured configs 2)Add remaining devices which don't have any configs stored		
Recovery:	Removing the devices and adding the devices again to fabric in following sequence 1)First add brownfield devices which have preconfigured configs 2)Add remaining devices which don't have any configs stored		

Parent Defect ID:	EFA-9456	Issue ID:	EFA-9456
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.3
Symptom:	Issue is seen when the devices which are being added to fabric have IP addresses already configured on interfaces and are conflicting with what EFA assigns.		
Condition:	Issue will be observed if devices being added to fabric have IP addresses assigned on interfaces and these IP addresses are already reserved by EFA for other devices.		

Parent Defect ID:	EFA-9456	Issue ID:	EFA-9456
Workaround:	Delete the IP addresses on interfaces of devices having conflicting configuration so that new IP addresses can be reserved for these devices. One way to clear the device configuration is using below commands: 1. Register the device with inventory efa inventory device register --ip <ip1, ip2> --username admin --password password 2. Issue debug clear "efa fabric debug clear-config --device <ip1, ip2>"		
Recovery:	Delete the IP addresses on interfaces of devices having conflicting configuration so that new IP addresses can be reserved for these devices. One way to clear the device configuration is using below commands: 1. Register the device with inventory efa inventory device register --ip <ip1, ip2> --username admin --password password 2. Issue debug clear "efa fabric debug clear-config --device <ip1, ip2>" 3. Add the devices to fabric		

Parent Defect ID:	EFA-9439	Issue ID:	EFA-9439
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	Dev-State and App-State of EPG Networks are not-provisioned and cfg-ready		
Condition:	Below are the steps to reproduce the issue: 1) Create VRF with local-asn 2) Create EPG using the VRF created in step 1 3) Take one of the SLX devices to administratively down state 4) Perform VRF Update "local-asn-add" to different local-asn than the one configured during step 1 5) Perform VRF Update "local-asn-add" to the same local-asn that is configured during step 1 6) Admin up the SLX device which was made administratively down in step 3 and wait for DRC to complete		
Workaround:	No workaround as such.		
Recovery:	Following are the steps to recover: 1) Log in to SLX device which was made admin down and then up 2) Introduce local-asn configuration drift under "router bgp address-family ipv4 unicast" for the VRF 3) Execute DRC for the device		

Parent Defect ID:	EFA-9398	Issue ID:	EFA-9398
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.3
Symptom:	Upon failed EFA upgrade, unable to revert to previous version of EFA.		

Parent Defect ID:	EFA-9398	Issue ID:	EFA-9398
Condition:	Upon failed EFA upgrade, unable to revert to previous version of EFA.		
Workaround:	Users can select "No" and re-try the operation using `efa deploy` on SLX or `source deployment.sh` from a server based deployment.		

Parent Defect ID:	EFA-9346	Issue ID:	EFA-9346
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.5.0
Symptom:	When the MCT Port channel membership is changed via config changes on SLX (out of band), Fabric service will not mark the device to be in cfg-refreshed state		
Condition:	This gives the user an incorrect impression that the status of the device is cfg-in-sync.		
Workaround:	Users can run manual Drift/reconcile command which will identify the drift and fix it. This condition will happen only when PO membership has changed out of band.		

Parent Defect ID:	EFA-9288	Issue ID:	EFA-9288
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.4
Symptom:	When device is removed from a non-clos fabric with multiple racks, BGP configuration is not removed.		
Condition:	Issue is seen ONLY in multi-rack non-CLOS fabric. When device is removed from a non-clos fabric with multiple racks, BGP configuration is not removed.		
Workaround:	Remove the configuration explicitly from device that is being removed from fabric by issuing "no router bgp" on SLX device.		
Recovery:	There is no recovery required as the device is decommissioned/removed from fabric. To remove stale BGP configuration from device user can issue "no router bgp" on the device directly.		

Parent Defect ID:	EFA-9065	Issue ID:	EFA-9065
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.3
Symptom:	EFA Port Channel remains in cfg-refreshed state when the port-channel is created immediately followed by the EPG create using that port-channel		

Parent Defect ID:	EFA-9065	Issue ID:	EFA-9065
Condition:	Below are the steps to reproduce the issue: 1. Create port-channel po1 under the ownership of tenant1 2. Create endpoint group with po1 under the ownership of tenant1 3. After step 2 begins and before step 2 completes, the raslog event w.r.t. step 1 i.e. port-channel creation is received. This Ralsog event is processed after step 2 is completed		
Recovery:	1. Introduce switchport or switchport-mode drift on the SLX for the port-channel which is in cfg-refreshed state 2. Perform manual DRC to bring back the cfg-refreshed port-channel back to cfg-in-sync		

Parent Defect ID:	EFA-9010	Issue ID:	EFA-9010
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.2
Symptom:	Creation of 100 Openstack VM/stacks fails at the rate of 10 stacks/min One stack has 1 VM , 2 networks and 3 Ports (2 dhcp and one nova port)		
Condition:	100 openstack stacks created at the rate of 10 stacks/min are sent to the EFA. The EFA processing requests at such high case resuts in overwhelming the CPU, Since the EFA cannot handle requests at such high rates, backlog of requests are created. This eventually results in VM reschedules and failure to complete some stacks with errors.		
Workaround:	100 openstack stacks can be created with lower rate of creation consistently eg 3 stacks/min		
Recovery:	Delete the failed or all openstack stacks and create them at lower rate of creation e.g 3 stacks/min		

Parent Defect ID:	EFA-8904	Issue ID:	EFA-8904
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.2
Symptom:	Single node deployment fails with 'DNS resolution failed.'		
Condition:	After a multi-node deployment and then un-deployment is done on a server, if single-node deployment is tried on the same server, the installer exits with the error, 'DNS resolution failed.'		
Workaround:	After un-deployment of the multi-node installation, perform a reboot of the server/ TPVM.		

Parent Defect ID:	EFA-8535	Issue ID:	EFA-8535
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	On a single-node installation of TPVM, after ip-change, EFA is not operational.		

Parent Defect ID:	EFA-8535	Issue ID:	EFA-8535
Condition:	After IP change of the host system, if 'efa-change-ip' script is run by a different user other than the installation user, in that case, EFA is not operational.		
Workaround:	Restart k3s service using the command 'sudo systemctl restart k3s'		

Parent Defect ID:	EFA-8448	Issue ID:	EFA-8448
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	When the ports provided by the user in “tenant update port-delete operation” contains all the ports owned by the port-channel, the PO goes into delete pending state. However, the ports are not deleted from the PO. They get deleted from the tenant though.		
Condition:	This issue is seen when the ports provided by the user in “tenant update port-delete operation” contains all the ports owned by the port-channel resulting in an empty PO.		
Workaround:	User needs to provide ports for “tenant update port-delete operation” which do not result in an empty PO i.e. PO needs to have at least 1 member port.		
Recovery:	Add the ports back using "tenant port-add operation" so that the port-channel has at least 1 member port. The use "efa configure tenant port-channel" to bring the po back to stable state.		

Parent Defect ID:	EFA-8297	Issue ID:	EFA-8297
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG update anycast-ip-delete operation succeeded for deletion of provisioned anycast-ip for admin-down device. This issue is observed only if an update anycast-ip-add operation is performed after device is put in admin down state and the new config is in non-provisioned state followed by anycast-ip-delete operation for already configured anycast-ip.		
Condition:	Steps to reproduce issue: 1) Configure EPG with anycast-ip (ipv4/ipv6) 2) Make one device admin-down 3) Anycast-ip update-add new anycast-ip (ipv6/ipv4) 4) Update-delete provisioned anycast-ip configured in step-1 (ipv4/ipv6) Step (4) should fail as IP is already configured on the device and trying to delete it should fail as part of APS.		
Workaround:	No workaround for this.		
Recovery:	Recovery can be done by configuring EPG again with the required configuration using efa or cleaning device config for anycast-ip on the switch.		

Parent Defect ID:	EFA-5928	Issue ID:	EFA-5928
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Configuring devices to default startup-config and adding them to a non-clos fabric does not enable all MCT ports resulting into fabric validation failure for missing link		
Condition:	Added devices immediately after setting to default startup config		
Workaround:	Remove the devices from fabric and re-add efa fabric device remove --name <fabric-name> --ip <device-ips> efa inventory device delete --ip <device-ips> efa fabric device add-bulk --name <fabric-name> --rack <rack-name> --username <username> --password <password> --ip <device-ips>		

Parent Defect ID:	EFA-5592	Issue ID:	EFA-5592
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Certificates need to be manually imported on replaced equipment in-order to perform RMA.		
Condition:	RMA/replaced equipment will not have ssh key and auth certificate, in-order to replay the configuration on new switch user needs to import the certificates manually.		
Workaround:	import certificate manually efa certificates device install --ips x,y --certType		