Extreme Fabric Automation, 2.3.1

Release Notes

9036773-01 Rev AB
November 2020

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/

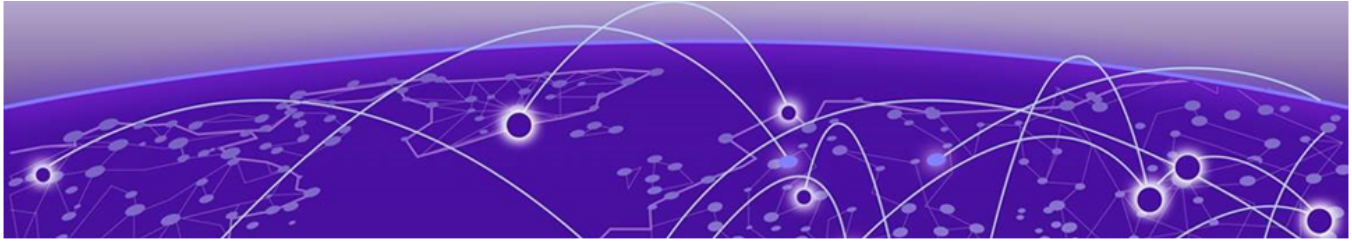

Table of Contents

Preface

This section describes the text conventions used in this document, where you can find additional information, and how you can provide feedback to us.

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.

Table 1: Notes and warnings (continued)

Icon	Notice type	Alerts you to...
	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 appears 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 Service Notifications

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

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form (all fields are required).

3. Select the products for which you would like to receive notifications.

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

4. Select **Submit**.

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

What's New in this Release

The following items are new or updated for Extreme Fabric Automation 2.3.1.

Table 4: Summary of changes

Change	Described in
Support for SLX 9250 as a border leaf Updated TPVM software support Updated TPVM deployment models Updated Neutron and SLX-OS compatibility	Supported Platforms and Deployment Models on page 9
New API to determine the status (up or down) of EFA: `https://<EFA-IP>:8078/v1/monitor/status/efa`	Monitor Service API Guide
Resolved and open defects The Open Defects list includes EFA-7446, which describes a workaround if EFA services are not running after you install the OVA.	Resolved Defects on page 11 and Open Defects on page 15

Supported Platforms and Deployment Models

Table 5: Bare Metal Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirement
EFA 2.1.0	External server (bare metal)	More than 24	Yes	16.04	CPU: 4 cores Storage: 50 GB RAM: 8 GB
EFA 2.2.0 and later	External server (bare metal)	More than 24	Yes	16.04, 18.04	

Table 6: OVA Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirement
EFA 2.1.0	External server (OVA)	More than 24	Yes	16.04	CPU: 4 cores Storage: 50 GB RAM: 8 GB
EFA 2.2.0 and later, in secure mode	External server (OVA)	More than 24	Yes	18.04	

Table 7: TPVM Deployment Models

EFA Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS version
2.1.0	SLX 9150 TPVM SLX 9250 TPVM	Up to 24	Yes	16.04	20.1.1
2.2.0	SLX 9150 TPVM SLX 9250 TPVM	Up to 24	Yes	18.04	20.1.x

Table 7: TPVM Deployment Models (continued)

EFA Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS version
2.3.0, 2.3.1	SLX 9150 TPVM SLX 9250 TPVM	Up to 24	Yes	18.04	20.1.2x
2.3.0, 2.3.1	SLX 9740 TPVM	Up to 24	Yes	18.04	20.2.2

Table 8: TPVM Software Support

TPVM Version	SLX-OS 20.1.1	SLX-OS 20.1.2x	SLX-OS 20.2.1x	SLX-OS 20.2.2	Ubuntu Version	EFA Version
3.0	Yes	Yes	No	No	16.04	2.1.0
4.0.0-5	No	Yes	No	No	18.04	2.2.0
4.0.0-7, 4.0.1	No	Yes	Yes	No	18.04	2.2.0, 2.3.0, 2.3.1
4.1.1-0	No	No	No	Yes	18.04	2.3.0, 2.3.1

Table 9: IP Fabric Topology Matrix

Platforms	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9030	18x.1.00b	✔				
SLX 9140	18s.1.03	✔				✔
SLX 9240	18s.1.03	✔	✔	✔		
SLX 9150	20.1.x, 20.2.x	✔				✔
SLX 9250	20.1.x, 20.2.x	✔	✔		✔	✔
SLX 9540	20.1.x, 20.2.x	✔			✔	
SLX 9640	20.1.x, 20.2.x				✔	
SLX 9740	20.2.x		✔	✔	✔	✔
SLX 9850	18r.1.00c		✔	✔		

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	2.3.1_02	20.1.2e

Resolved Defects

Parent Defect ID:	EFA-7341	Issue ID:	EFA-7341
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	"efactl status" or any "k3s kubectl" commands respond with the following error message: The connection to the server 127.0.0.1:6443 was refused - did you specify the right host or port? [Ref: EMAX-105]		
Condition:	The k3s datastore cluster may fail to restore when the standby node's management interface is shut down for sometime and brought back up again. efa-monitor incorrectly recovered the standby node as a new standalone cluster. Once the management interface was restored, the active node tries to join this new cluster and fails to join because the transactions are ahead of the standby node. The standby node should have remained down in a fault state and joined the current active once the management interface was restored.		
Workaround:	The issue has been fixed in 2.3.1		
Recovery:	The issue has been fixed in 2.3.1		

Parent Defect ID:	EFA-7346	Issue ID:	EFA-7346
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Except local-ip other configs are not removed from the "Admin-UP" device after EPG delete operation on partial success topology [Ref: EMAX-106]		
Condition:	This issue can be observed if EPG deletion containing local-ips has failed because any one of the devices is admin down.		
Workaround:	The issue has been fixed in 2.3.1		
Recovery:	The issue has been fixed in 2.3.1		

Parent Defect ID:	EFA-7351	Issue ID:	EFA-7351
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Deletion of the EPG fails with an error " Error: Operation "epg delete" not allowed on an "epg in uninitialised state".		
Condition:	Below is the scenario in which the issue can happen: 1.EPG creation is in progress. 2.Tenant service restarts while (1) is in progress. 3.Deletion of EPG (created in step 1) is attempted after the tenant service restart.		
Workaround:	The issue has been fixed in 2.3.1		
Recovery:	The issue has been fixed in 2.3.1		

	EFA-7355	Issue ID:	EFA-7355

Parent Defect ID:			
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Even though the configuration is in sync between the EFA and SLX, the switch doesn't have: 1. "cluster-track" configuration under the CEP port-channel. 2. "graceful-restart" configuration under the router bgp address-family.		
Condition:	The issue can occur in the below scenario: 1. EFA 2.1.0 version is installed. 2. CEP port-channels are configured. 3. L3 EPGs are configured with the CEP port-channel and anycast-ipv4. 4. EFA 2.1.0 is upgraded to EFA 2.2.0. 5. efa-db-upgrade-from-2-1-0.sh is executed. 6. EFA 2.2.0 is upgraded to EFA 2.3.0. 7. Manual drift and reconcile is executed using "efa inventory drift-reconcile execute --ip <switch-ip> --reconcile" CLI.		
Recovery:	A temporary EPG can be created using the "CEP port-channel" and the "VRF" for which the "cluster-track" and "graceful-restart" configuration was missing respectively on the switch, so that the missing "cluster-track" and "graceful-restart" configurations get pushed to the switch during the EPG create.		

Parent Defect ID:	EFA-7357	Issue ID:	EFA-7357
Severity:	S3 – Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	3DES ciphers from Golang TCP/8078 HTTPS server		
Workaround:	As part of Security hardening process this issue has been fixed in EFA 2.3.1		
Recovery:	As part of Security hardening process this issue has been fixed in EFA 2.3.1		

Parent Defect ID:	EFA-7358	Issue ID:	EFA-7358
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	The raslog TCP/6514 on EFA-2.3 offers both TLS 1.0 and TLS 1.1		
Workaround:	This issue is fixed in EFA 2.3.1 by way stricter enforcement TLS Version => 1.2		
Recovery:	This issue is fixed in EFA 2.3.1		

Parent Defect ID:	EFA-7360	Issue ID:	EFA-7360
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	system set-mtu and set-admin-state commands fail due to 'Service Unavailable' error [Ref: EMAX-107]		
Condition:	when the Keystore value of the device are not available in the Asset Services.		

Workaround:	This issue is fixed in EFA 2.3.1
Recovery:	This issue is fixed in EFA 2.3.1

Parent Defect ID:	EFA-7362	**Issue ID:**	EFA-7362
Severity:	S2 – High		
Product:	Extreme Fabric Automation	**Reported in Release:**	EFA 2.3.0
Symptom:	The creation of 10 MS networks with subnets parallelly using hot template is failing [Ref: EMAX-108]		
Condition:	Creation of ports aross 10 networks on higher number of segments takes more time leading to failure.		
Workaround:	Issue is fixed in 2.3.1		

Parent Defect ID:	EFA-7364	**Issue ID:**	EFA-7364
Severity:	S3 – Medium		
Product:	Extreme Fabric Automation	**Reported in Release:**	EFA 2.3.0
Symptom:	The pod status shown in efactl status for a given node as it is coming up after an EFA HA failover makes it difficult to determine if the node is in active/standby state to be ready to service EFA requests (Ref: EMAX-109)		
Condition:	This is seen right after an EFA HA failover as the nodes are transitioning from active/ standby state.		
Workaround:	This issue is fixed in 2.3.1 adding an API to correctly determine each node's status. An example of using the API is as follows curl --location --request GET '<monitor-endpoint>/v1/monitor/status/efa'		
Recovery:	This issue is fixed in 2.3.1 adding an API to correctly determine each node's status. An example of using the API is as follows curl --location --request GET '<monitor-endpoint>/v1/monitor/status/efa'		

Parent Defect ID:	EFA-7365	**Issue ID:**	EFA-7365
Severity:	S3 – Medium		
Product:	Extreme Fabric Automation	**Reported in Release:**	EFA 2.3.0
Symptom:	The inability to run an upgrade process until the underlying prerequisite check is fixed. efa login and all other efa commands fail to execute, showing this message in efa login `CLI is not registered as a client. Please run 'source /etc/profile' to update your environment.` [Ref: EMAX-110]		
Condition:	The issue is seen when there is a failing prerequisite check at the start of the upgrade process. After the failure is reported and the upgrade process is cancelled without undeploying the failed upgrade, efa login and thus all other efa commands fail to run.		
Workaround:	This issue is fixed in EFA 2.3.1		
Recovery:	This issue is fixed in EFA 2.3.1		

	EFA-7372	**Issue ID:**	EFA-7372

Parent Defect ID:			
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Unable to reach VMs (ping) through new DHCP agent after compute poweroff [Ref: EMAX-111]		
Condition:	When a DHCP agent is re-located to other compute host, change in the switch details cause port configuration to fail.		
Workaround:	This Issue is fixed in EFA 2.3.1		
Recovery:	This Issue is fixed in EFA 2.3.1		

Parent Defect ID:	EFA-7386	Issue ID:	EFA-7386
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	In a multinode EFA HA deployment the following symptoms are noticed for this issue after an HA failover scenario: efa login fails and efactl status shows one or both nodes as a not ready state. k3s systemctl service fails to start on the not ready node.		
Condition:	The issue is found to occur when the system is under stress due to prolonged failover testing, causing the maxium inotify watchers limit to be hit for the k3s process. This caused the k3s service to fail and hence make efa non functional.		
Workaround:	This issue is fixed in EFA 2.3.1		
Recovery:	This issue is fixed in EFA 2.3.1		

Parent Defect ID:	EFA-7388	Issue ID:	EFA-7388
Severity:	S2 – High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Duplicate networks with same vlan in efa is observed [Ref: EMAX-115]		
Condition:	“openstack network create” should have failed.		
Workaround:	Issue is fixed in 2.3.1		

Parent Defect ID:	EFA-7392	Issue ID:	EFA-7392
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	On an HA setup the monitoring service is not accessible with https protocol intermittently		
Condition:	After HA failover		
Workaround:	This issue is fixed in EFA 2.3.1		
Recovery:	This issue is fixed in EFA 2.3.1		

Open Defects

Parent Defect ID:	EFA-2857	Issue ID:	EFA-2857
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Addition of fabric ports and MCT ports to the tenant succeeds even though these ports cannot be used to connect to any endpoints.		
Condition:	Addition of fabric ports (ports connecting leaf to spine, spine to super-spine etc) and MCT ports (ports connecting the multi-homed leaf) to the tenant.		
Workaround:	Fabric ports and MCT ports to be not added to the tenant.		

Parent Defect ID:	EFA-3512	Issue ID:	EFA-3512
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG create might fail with VNI resource not being available in the fabric		
Condition:	Execute EPG create, delete and re-create CLI in quick succession as below: 1. Create EPG/Networks with user-provided VNI parameter. 2. Delete EPG. 3. Create EPG again with the same parameters as in step-1.		
Workaround:	Provide a wait of 30 seconds between the create and delete CLI on the same EPG.		

Parent Defect ID:	EFA-3717	Issue ID:	EFA-3717
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.1.0
Symptom:	BD based L3 EPG networks will go into cfg-refresh-err state if out-of-band configured BD-LIFs (without associated VLAN) are present on the switch.		
Condition:	1. Create BD-LIF (without any associated VLAN) configuration manually on the switch followed by EFA inventory update. 2. Create EPG which results in the same BD-LIF on the switch.		
Workaround:	Cleanup the out-of-band created LIFs from the switch before EPG create,		
Recovery:	1. Delete the EPG which is in cfg-refresh-err state 2. Clean up the out-of-baned created LIFs from the switch. 3. EFA inventory update to reflect the modified config from the switch. 4. Create the EPG again.		

Parent Defect ID:	EFA-4056	Issue ID:	EFA-4056
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0

Symptom:	The second EPG will get created successfully but the anycast-ipv6 provided in the second EPG will not get configured on the switch. The second EPG create command should have failed with an error mentioning a unique anycast-ipv6 needs to be provided for a given BD name.
Condition:	1. Create BD based L3 EPG with some BD label and an anycast-ipv6. 2. Create another BD based L3 EPG with the same BD label (as used in 1) but with a different anycast-ipv6.
Workaround:	Provide a unique anycast-ipv6 for a given BD name across the L3 EPG during EPG create and update operations.
Recovery:	Remove the subsequent input incorrect anycast-ipv6 from the EPG and provide a unique anycast-ipv6 for a given BD name across the L3 EPG during L3 EPG create and update operations.

Parent Defect ID:	EFA-4832	Issue ID:	EFA-4832
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	On any operation on EFA services, the error 'Invalid credentials for device 10.x.x.x.' is displayed.		
Condition:	When a device is registered with device credentials in EFA and the same credentials are deleted or updated in SLX.		
Workaround:	Any update in device credentials on SLX(deleting the user, changing password etc) has to be updated in EFA using EFA CLI.		

Parent Defect ID:	EFA-5064	Issue ID:	EFA-5064
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EFA configures both "suppress-arp" (needed for ipv4) and "suppress-nd" (needed for ipv6) for the particular network (VLAN/BD). EFA shouldn't have configured "suppress-nd" on the network.		
Condition:	Configure L3 EPG with only anycast-ip and no anycast-ipv6.		
Recovery:	No recovery needed.		

Parent Defect ID:	EFA-5257	Issue ID:	EFA-5257
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	When VRF is added and deleted to/from and Endpoint Group, in quick succession, multiple times, events received from inventory service can get interleaved with the commands. This can cause EFA command execution path to find database entries that are yet to be deleted due to previous command run.		
Condition:	Issue can be observed when vrf-add and vrf-delete operation is executed on Endpoint Group in quick succession multiple times.		

Workaround:	Workaround is to wait for a few minutes before executing the vrf-add again on Endpoint Group
Recovery:	No recovery required

Parent Defect ID:	EFA-5286	Issue ID:	EFA-5286
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Responses to REST requests for VRF Update do not contain details of specific errors that occurred during rollback of errored configuration.		
Condition:	Issue is observed when VRF parameter update is requested, error occurs and rollback that is triggered also encounters error.		
Workaround:	The final error encountered is visible as part of the final error returned to the REST request, in string form.		
Recovery:	No recovery required		

Parent Defect ID:	EFA-5287	Issue ID:	EFA-5287
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Drifted configuration under VRF is not reflected in ‘efa tenant vrf show’ output. The VRF is shown as being in ‘cfg-in-sync’ state.		
Condition:	Issue is observed when configuration under VRF has drifted on the switch due to various reasons.		
Workaround:	Workaround is to use display of ‘efa tenant epg show’ to determine if there is a drift in configuration.		
Recovery:	No recovery required		

Parent Defect ID:	EFA-5590	Issue ID:	EFA-5590
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EFA CLI behavior on pressing <tab> following a complete command will execute the command handler or will get stuck in cli if waiting on user input on background		
Condition:	(efa:root)root@ubuntu:~/efa/efa_40# efa inventory rma execute --ip 10.24.95.157 --co<tab> - cli gets stuck and needs ^C^C to exit, as it is waiting on user input on backend (efa:extreme)extreme@tpvm2:~$ efa fabric create --name fab4 – <tab> - this creates a fabric		

Workaround:	These can be avoided by using '?' before the tab and using the complete command options. In most cases this behavior will not cause an issue, but issue may be seen in cases where there are multiple optional keywords eg in fabric-create, rma execute, fabric add bulk , etc.
Recovery:	Press ^C^C to exit a command thats stuck on cli. Reverse any command executed as above, if it is not intended.

Parent Defect ID:	EFA-5675	Issue ID:	EFA-5675
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG update "vrf-add" operation can fail with the reason as vrf to be added has conflicting vrf on the switch.		
Condition:	Execute EPG update "vrf-add", "vrf-delete" and "vrf-add" operation CLI in quick succession as below: 1. Update EPG for operation vrf-add. 2. Update EPG for operation vrf-delete. 3. Update the same EPG again with operation vrf-add for the same vrf which was deleted in step 2.		
Workaround:	Provide a wait of 30 seconds between the EPG update vrf-add and vrf-delete operations on the same EPG.		
Recovery:	-		

Parent Defect ID:	EFA-5689	Issue ID:	EFA-5689
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	VRF Backup routing configuration on the SLXOS-9140 will be inadequate and hence the backup routing functionality will not work on SLXOS-9140.		
Condition:	L3 EPG create or update operation with the member ports residing on SLXOS-9140.		
Workaround:	Backup routing needs to be disabled at the fabric setting level if the fabric has SLXOS-9140 devices.		
Recovery:	Backup routing needs to be disabled at the fabric setting level if the fabric has SLXOS-9140 devices.		

Parent Defect ID:	EFA-5708	Issue ID:	EFA-5708
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Device status remains in cfg-in-sync when devices have mismatch in selective configs w.r.t to intended configs		
Condition:	Periodic discovery identifies the drift and raise appropriate events; most of the events that are handled do not change the device state from cfg-in-sync to cfg-refresshed		

Workaround:	Drift and Reconcile events fix the problem
Recovery:	Execute drift and reconcile

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-5821	Issue ID:	EFA-5821
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG create fails after a wait of more than 16 minutes.		
Condition:	Create EPG having more than 256 networks.		
Workaround:	Incrementally add networks to the EPG. Example as below: efa tenant epg create --tenant <> --name <> --ctag-range 2-256 efa tenant epg update --tenant <> --name <> --operation ctag-range-add --ctag-range 257-512		
Recovery:	Verify the EPG/Networks state using "efa tenant epg show" after CLI completes and correct the EPG/Networks state using "efa tenant epg update" CLI.		

Parent Defect ID:	EFA-5822	Issue ID:	EFA-5822
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG delete fails after a wait of more than 16 minutes.		
Condition:	Delete EPG having more than 256 networks.		
Workaround:	Incrementally delete networks from the EPG and finally delete EPG Example as below: efa tenant epg update --tenant <> --name <> --operation ctag-range-delete --ctag-range 1001-1256 efa tenant epg delete --tenant <> --name <>		
Recovery:	Verify the EPG is deleted using "efa tenant epg show" after CLI completes.		

	EFA-5834	Issue ID:	EFA-5834

Parent Defect ID:			
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG creation with 4K ctags can result in switch going un-responsive		
Condition:	EPG create operation with 4K ctags.		
Workaround:	Incrementally add networks to the EPG. Example as below: efa tenant epg create --tenant <> --name <> --ctag-range 2-1000 efa tenant epg update --tenant <> --name <> --operation ctag-range-add --ctag-range 1001-2000		

Parent Defect ID:	EFA-5841	Issue ID:	EFA-5841
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-5927	Issue ID:	EFA-5927
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Configuration reconciliation fails with an error "drift and reconcile failed waiting for status from tenant." because of the timeout.		
Condition:	When the switch configurations drift from the intended configurations in EFA due to scenarios as follows: 1. L3 Epg is created with large ctag-range (e.g. 2-2000) 2. EFA configured VLANs and PO configurations are manually removed from the switch. 3. Switch is reloaded in maintenance mode		
Recovery:	After the switch is moved out of maintenance mode after reload, configuration drift can be viewed and reconciled using "efa inventory drift-reconcile execute --reconcile --ip <switch-ip>" CLI.		

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-5936	Issue ID:	EFA-5936
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Overlay-Gateway/Cluster Instance is remaining on the switch.		
Condition:	Configure EPG/Networks on a fabric device and delete device from the fabric.		
Recovery:	Manually delete the stale Overlay-Gateway and Cluster instances from the switch.		

Parent Defect ID:	EFA-5952	Issue ID:	EFA-5952
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	While executing drift-reconcile port-channels created between MCT pairs and are part of EPG's will be shown as drifted even though the SLX running config is in sync.		
Condition:	This happens when drift is executed while SLX is in maintenance mode.		
Recovery:	Turn off maintenance mode and execute inventory device update which will sync the right status.		

Parent Defect ID:	EFA-6318	Issue ID:	EFA-6318
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.1
Symptom:	In Case of Clos Fabric , Cluster drift gets identified for all the devices which are part fabric as soon as fabric gets created even though actual drift is not present on the SLX switch.		
Condition:	This issue observed with below steps, 1)Created CLOS Fabric with multiple MCT-Pairs and a Spine device 2)When checking drift for all devices part of fabric, Cluster Config can be seen as refreshed, but no impact can be seen on fabric show. Also drift status for fabric states false in drift o/p.		
Workaround:	EFA and SLX will be in sync when the drift and reconcile (triggered because of maintenance mode enable-on-reboot) 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. - Create VRF/EPG having route target, static route and bgp configuration. - Introduce drift in VRF route target or static route or bgp configuration on SLX switch. - Update device from efa command "efa inventory device update --ip <device ip>" - Check device drift using efa command as "efa inventory drift-reconcile execute --ip <device ip>" - 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-6832	Issue ID:	EFA-6832
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	If any of the delete nature CLI of EPG/PO has failed then that EPG/PO will be stuck in delete-pending state and no other operation will be allowed on that EPG/PO except delete.		
Condition:	This behaviour is correct if delete operation has failed because of device connectivity or device de-configuration issues. But if it is failed because validations then EPG/PO should remain in their original state and not in the delete-pending state.		
Workaround:	None		
Recovery:	There are 2 recovery options: 1. Fix the reason for failure on the device and re-run the same CLI which will complete the operation and EPG/PO will back to the there expected state. 2. Use --force operation if EPG/PO needs to be completely deleted from EFA.		

Parent Defect ID:	EFA-6835	Issue ID:	EFA-6835
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	EFA CLIs did not respond following bulk add of devices		
Condition:	Register more 7 devices into EFA at same time using add-bulk fabric CLI in HA multi-node setup.		
Workaround:	Register less then 5 devices in EFA at same time.		
Recovery:	Restart the EFA sever by reloading the Switch.		

Parent Defect ID:	EFA-7062	Issue ID:	EFA-7062
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Device settings for Maintenance mode in EFA is not updated when the device's state changes.		

Condition:	Device may come out of Maintenance mode state based on actions from EFA like Drift-reconcile or a change via CLI in SLX device. The latest state of device will not be automatically updated in the 'efa inventory device settings show' in efa
Workaround:	Re-issue the 'efa inventory device settings update --maint-mode-enable yes ' command to execute the required setting or validate the setting on SLX device using 'show system maintenance'
Recovery:	Re-issue the 'efa inventory device settings update --maint-mode-enable yes ' command to execute the required setting or validate the setting on SLX device using 'show system maintenance'

Parent Defect ID:	EFA-7080	Issue ID:	EFA-7080
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Tenant commands can take more time (> the timeout =1000 seconds) to complete and hence the failure. The status of timed out fail command can be seen using the relevant "show" command.		
Condition:	The issue can happen in the below scenario: 1. EFA has a scaled tenant configuration. 2. Significant configuration drift introduced between SLX and EFA. 3. Inventory device update is performed resulting in configuration drift derivation of the tenant entities. 4. Tenant commands are executed when (3) is in progress.		
Workaround:	Tenant commands should be attempted when the system is in a stable state i.e. when the EFA tenant is not busy deriving the scaled configuration drift of the tenant entities.		

Parent Defect ID:	EFA-7114	Issue ID:	EFA-7114
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Drift is shown for "pseudowire-profile" configuration in command, "efa inventory drift-reconcile execute --ip <device ip>"		
Condition:	This is observed with below steps - Create fabric with back routing enable and configure the fabric. - Create epg and delete the same epg. - Make one of cluster pair device as admin down. - Create same epg which was created and deleted above. - Make the admin down device as admin up. - "pseudowire-profile" still remain as drifted.		
Workaround:	None		
Recovery:	To recover please follow below steps, - Delete "pseudowire-profile" from SLX switch. - Execute drift reconcile as "efa inventory drift-reconcile execute --ip <SLX switch IP> --reconcile"		

	EFA-7119	Issue ID:	EFA-7119

Parent Defect ID:			
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	If any delete nature CLI of "service bgp peer" has failed then bgp service status is not moving to the "delete-pending" state.		
Condition:	Any delete nature CLI of "service BGP peer" has failed because any of the below-listed reasons will lead the BGP service to this state. 1. Device connectivity issues 2. Failed to deconfigure the peer on device 3. Targeted topology falls into the APS category.		
Workaround:	None		
Recovery:	There are no side effects of this issue except "update" operations on that BGP service will be allowed which is not expected behavior after the "service bgp peer delete" operation has failed. There are 2 recovery options: 1. Fix the reason for failure and again try to delete BGP service. 2. Delete BGP service using "--force" option. In this case, configs will remain on admin down the device.		

Parent Defect ID:	EFA-7132	Issue ID:	EFA-7132
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Vlan description is missing after reconciliation		
Condition:	- Remove the VLAN configuration from the device - Reload the device with "maintenance mode on reboot enabled" - Description configuration is missing for reconciled VLANs		
Recovery:	Add VLAN description on the device		

Parent Defect ID:	EFA-7142	Issue ID:	EFA-7142
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Pressing the tab key in the EFA command line when the current command is complete and has no further available options can result in the command being executed silently by the shell.		
Condition:	This is not reliably reproducible.		
Workaround:	EFA CLI completion can be disabled. This does not affect EFA CLI operation or limit the commands available; it prevents convenient discovery of extra sub-commands and flags by using the tab key from within the bash shell. To disable, do the following: - remove /etc/bash_completion.d/efa - comment out or remove the line "source /etc/bash_completion.d/efa" from /etc/bash.bashrc - start a new shell for EFA CLI use		

	EFA-7144	Issue ID:	EFA-7144

Parent Defect ID:			
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	'Drift and Reconcile' for a device is skipped when EFA admin of device is performed with 'force' option		
Condition:	When 'efa inventory admin-state up' command is issued with a 'forcce' option		
Workaround:	The Drift and Reconcile can be run manually using the below command: 'efa inventory drift-reconcile execute '		

Parent Defect ID:	EFA-7155	Issue ID:	EFA-7155
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Executing the command “efa-sync execute” when one of the SLX switches in cluster is unreachable, could take several minutes to complete. This happens only when multiple neutron entities have to be synced with EFA.		
Condition:	The command takes longer time to complete only when: a)Out of sync configuration is large. That is, many neutron entities are out of sync on EFA. b)One of the SLX nodes is not reachable.		
Workaround:	No workaround available to make the operation faster on EFA		
Recovery:	efa-sync execute becomes faster after the SLX nodes becomes reachable		

Parent Defect ID:	EFA-7159	Issue ID:	EFA-7159
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	MCT Port Channel Admin State is not reconciled		
Condition:	Issue will be observed when 1. MCT PO connecting the leaf pair is administratively shut-down explicitly via out of band management application like CLI and 2. maintenance-mode is set to enable-on-reboot and 3. Device is reloaded. After reload, Drift and reconcile doesn't set the admin state of MCT PO to UP.		
Workaround:	Workaround: 1. After the device is reloaded and switch is taken out of maintenance mode, issue a "no shutdown" on the MCT PO interface on the device either through CLI or EFA exec mode. 2. Issue a manual drift and reconcile on the devices where this MCT PO is configured.		
Recovery:	No recovery required		

Parent Defect ID:	EFA-7164	Issue ID:	EFA-7164
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0

Symptom:	For Static Route/Static Route-BFD delete operation, if deletion has failed for any device then configs will be deleted from all other devices but they will remain in EFA for all the devices, and "efa tenant vrf show" CLI will continue to show all of them to the user and VRF will remain in the delete-pending state.
Condition:	SR/SR-BFD delete operation in vrf update has failed because of bellow conditions 1. device connectivity issues 2. Failed to deconfigure the peer on device 3. Targeted topology falls into the APS category.
Workaround:	None
Recovery:	Delete SR/SR-BFD using the same "VRF update" command after fixing the issue which was causing the failer.

Parent Defect ID:	EFA-7181	Issue ID:	EFA-7181
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Reconcilation of drifted Tenant Configuration fails for scaled configuration		
Condition:	Reconcile fails for a device on which EFA has configured many L3 networks (more than 300+ Networks) by one or more tenants.		
Recovery:	To reconcile the device having 300+ L3 Networks, first replay the last saved good configuration and execute reconcile command for the device. 1. Replay Config: efa inventory config-replay execute --ip <device-ip> --uuid <saved-config> 2. Reconcile: efa inventory drift-reconcile execute --ip <device-ip>		

Parent Defect ID:	EFA-7200	Issue ID:	EFA-7200
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	"router bgp" configs will be missing on the devices but Tenant service will not show drift for those configs and hence it will not reconcile those missing configs.		
Condition:	This is observed with the below steps 1. Put SLXs in maintenance mode 2. Remove the router bgp configuration, few VRFs from the device. 3. Reboots both the MCT nodes where TPVMs containing EFA are hosted. 4. Verify after 30 mins, the switches are out of maintenance mode. 5. After a few iterations most of the BGP configs are missing.		
Workaround:	None		
Recovery:	Delete missing configs from EFA also and recreate them.		

Parent Defect ID:	EFA-7234	Issue ID:	EFA-7234
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Reconciliation of drifts in the port-channel and description does not happen		

Condition:	This problem occurs only when drift is introduced specifically on speed and description along with admin-state or cluster-client configuration.
Workaround:	None
Recovery:	Remove the port-channel from device and execute manual reconcile.

Parent Defect ID:	EFA-7237	Issue ID:	EFA-7237
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	EFA system backup is generated using efa system backup command. Undeply EFA and install fresh EFA Restore EFA backup generated earlier using "efa system restore" command The restore process preserves the config backup interval and count. However the auto config backup process is not taking backups as configured.		
Condition:	Auto Config backup process does not take expected number of backups as configured.		
Workaround:	User can take config backup manually using "efa inventory config-backup execute".		
Recovery:	Fresh install or restart of EFA application would start the auto config backup process		

Parent Defect ID:	EFA-7261	Issue ID:	EFA-7261
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	"efa tenant epg show" and "efa tenant po show" will not reflect correct drift status if aggregator ports are removed.		
Condition:	This display issue is seen only when aggregator port is removed from device.		
Workaround:	None		
Recovery:	Manual execution of reconcile will sync the device and EFA database. This will in turn resolve the display issue.		

Parent Defect ID:	EFA-7264	Issue ID:	EFA-7264
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	After SLX Switch reload, Configuration drift in port channel members do not get reconciled by EFA		
Condition:	Below are the steps when the configuration drift created does not get reconciled by EFA: - Configuration of fabric/tenant/po/vrf/epg - 'maintenance mode enable-on-reboot ' enabled on the switch - Drift in configuration created by deleting port channel - Reloading of the switch - As part of reconcilation, port channel was created but Eth <port> was not aggregated to channel-group.		

Workaround:	None
Recovery:	Execute manual Drift and Reconcile using this command: " efa inventory drift-reconcile execute --ip <IP of device> --reconcile "

Parent Defect ID:	EFA-7269	Issue ID:	EFA-7269
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	During drift reconcile triggered because of maintenance mode enable-on-reboot configuration, interface configurations shown as drifted even though actual drift is not present on the SLX switch.		
Condition:	This issue observed with below steps, - Configured fabric/tenant/po/vrf/epg/bgp peer/peer-group - Enabled maintenance mode enable-on-reboot on SLX switch. - Reload SLX switch. - Drift and Reconcile process shows drift for interface used in EPG which was not drifted.		
Workaround:	None		
Recovery:	EFA and SLX will be in sync when the drift and reconcile (triggered because of maintenance mode enable-on-reboot) is completed.		

Parent Defect ID:	EFA-7304	Issue ID:	EFA-7304
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	While trying to upgrade from secure mode setup, the EFA installer prompts the user to choose between Secure mode and Standard mode.		
Condition:	This will be seen when proceeding through the menu options whenever an upgrade is performed.		
Workaround:	The user should choose secure if the current install is a secure-install. The user should choose non-secure if the current install is a non-secure install.		

Parent Defect ID:	EFA-7312	Issue ID:	EFA-7312
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Device setting parameters like health check enable, heartbeat miss threshold etc. shows as default values. Drift-reconcile get triggered after default heartbeat miss threshold.		

Condition:	HA double fault failure.
Recovery:	Update device setting parameters once again using CLI. $ efa inventory device setting update --ip <Device-ip> --maint-mode-enable-on-reboot <yes> --health-check-enable <yes> --health-check-interval <7m> --health-check-heartbeat-miss-threshold <3> --config-backup-periodic-enable <yes> --config-backup-interval <4m> --number-of-config-backups <10>

Parent Defect ID:	EFA-7322	Issue ID:	EFA-7322
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Firmware-download 'Update State' shows 'In Progress' indefinitely when triggered on Active Node's SLX		
Condition:	When SLX firmware upgrade is initiated from the EFA (deployed in HA mode) on the same SLX device where EFA is co-resident and is current active node.		
Workaround:	Recommended way to perform an SLX firmware upgrade from the EFA HA cluster being hosted by the same SLX devices: Upgrade the standby EFA node SLX device. Perform HA failover by stop/start active TPVM Upgrade the old active or new standby EFA node SLX device.		
Recovery:	workaround to allow the firmware download status details to be read from the device and updated in inventory DB. This can be triggered by the command - "efa inventory device update'		

Parent Defect ID:	EFA-7326	Issue ID:	EFA-7326
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	VRF used in EndpointGroup having invalid Target Network cause Drift & Reconcile operation to fail with error- "Target network is invalid"		
Condition:	1) VRF is created with Invalid Target Network in earlier version of EFA (before EFA 2.3.0) 2) EndpointGroup is created with same VRF in earlier version of EFA (before EFA 2.3.0) 3) Drift and Reconcile operation is performed		
Workaround:	NA		
Recovery:	1) Delete EndpointGroup using VRF having invalid Target Network 2) Delete VRF 3) Recreate VRF with valid Target Network 4) Recreate EndpointGroup with new VRF		

Parent Defect ID:	EFA-7345	Issue ID:	EFA-7345
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	A single-node EFA install at version 2.2 failed to upgrade to 2.3.		

Condition:	EFA 2.2 to EFA 2.3 upgrade fails only when EFA 2.2 was upgraded (not fresh install) from EFA 2.1.
Recovery:	Please see the section "Recover from an Upgrade Failure" in the EFA 2.3 Deployment Guide.

Parent Defect ID:	EFA-3694	Issue ID:	EFA-3694
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Request processing takes longer time and CLI remains hung till the request is processed. EPG update may fail if the wait time is more than 16 minutes.		
Condition:	1. Create EPG with 100 networks. 2. Update EPG with 20 ports/port-channels from one switch using "port-group-add" operation.		
Workaround:	There is no workaround.		

Parent Defect ID:	EFA-7389	Issue ID:	EFA-7389
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	In a very rare scenario, http daemon may not start		
Workaround:	Reload SLX device		

Parent Defect ID:	EFA-7396	Issue ID:	EFA-7396
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	If EPG is in delete-pending state with partially cleaned up vlan config on port-channels, deleting the EPG could fail. EPG could remain in delete-pending state unless deleted with force option.		
Condition:	This condition can happen if EPGs in delete-pending state have partially cleaned up vlan configuration on port-channels. That is, the vlan configuration on port-channel is cleaned up on some devices. Deleting the EPG errors out stating that the configuration does not exist.		
Workaround:	Delete the EPG using 'force' option.		
Recovery:	No recovery required.		

Parent Defect ID:	EFA-7398	Issue ID:	EFA-7398
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	Tenant drift reconcile operation failed with error message "Drift generation failed for the device"		

Condition:	This can happen when EPGs are created with different devices using same/common ctag/network and one or more of such common ctag/network has drifted on one of the device
Recovery:	1) Delete EndpointGroup(s) using same/common ctag/network(s) on the device where drift generation has failed 2) Let Drift-Reconcile operation complete for Tenant 3) Recreate deleted EndpointGroup(s)

Parent Defect ID:	EFA-7403	Issue ID:	EFA-7403
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	Devices failed to add into fabric		
Condition:	i) Delete device from fabric ii) Re-add same device into fabric right away.		
Workaround:	i) Wait for 4-5 min. ii) Re-add devices into fabric by re-execute same fabric bulk add CLI. [efa fabric device add-bulk --name small-fabric --border-leaf <device-ip> --username <name> --password <pass>]		
Recovery:	i) Wait for 4-5 min. ii) Re-add devices into fabric by re-execute same fabric bulk add CLI. [efa fabric device add-bulk --name small-fabric --border-leaf <device-ip> --username <name> --password <pass>]		

Parent Defect ID:	EFA-7404	Issue ID:	EFA-7404
Severity:	S1 - Critical		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	After firmware download devices from fabric shows a "cfg refresh error" state.		
Condition:	Inventory service discovery happened before lldp links got enabled on devices.		
Workaround:			
Recovery:	i) Wait for few 5-10 min. ii) Update devices into Inventory. [$ efa inventory device update --ip <IP address of devices separated by ','>] iii) Complete firmware download for remaining device into fabric. iv) Reconfigure the fabric to recover from "cfg resfreshed" state [$ efa fabric configure --name <Name>]		

Parent Defect ID:	EFA-7446	Issue ID:	EFA-7446
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	EFA services are not running after you install the efa-2.3.1.ova.		

Condition:	This is an inherent issue with the OVA.

Workaround:	1. Log in to the OVA.
	2. Confirm that the k3 service is running.
	`sudo service k3s status`
	`k3s.service - Lightweight Kubernetes ...`
	`Active: active (running) since Sat 2020-10-31 00:08:24 UTC; 2 days ago`
	3. Run the following script manually: `/opt/godcapp/efa/adjust_single_node_ip_change.sh`
	4. Reboot the EFA server.
	5. Log in to EFA.
	6. Verify that EFA services are running.
	`efactl status`