

February 2020



Extreme Fabric Automation 2.1.0 Release Notes

Part number: 9036615-00 Rev AA

© 2020, Extreme Networks, Inc. All Rights Reserved.

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

Contents

Document History	4
Preface	5
Introduction	6
Key Features.....	6
Software Support Matrix	7
EFA Core Services and Integrations	7
New Features	8
Prerequisites for SLX Devices.....	12
Prerequisites for the vCenter Controller and Compute Nodes	12
Prerequisites for the OpenStack Controller and Compute Nodes.....	12
Prerequisites for the OpenStack Controller and Compute Nodes.....	13
OpenStack EFA Neutron Plugin.....	13
Fresh Install of EFA on TPVM	13
Fresh Install of EFA on a Server.....	14
Fresh Install of EFA on a VM using OVA.....	14
Upgrade EFA on TPVM.....	14
Upgrade EFA on a Server	15
EFA Database Backup and Restore	16
Known Limitations and Issues.....	16
EFA Documentation	18

Document History

Version	Summary of Changes	Publication Date
1.0	Initial Release for 2.1.0	February 2020

Preface

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, call (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 want to receive notifications.
Note: You can change your product selections or unsubscribe at any time.
4. Select **Submit**.

Extreme Resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at www.extremenetworks.com/support/documentation.

Introduction

Extreme Fabric Automation (EFA) is a microservices-based application that manages the life cycle of IP Fabric CLOS and Small Data Center deployments. All of the microservices support REST APIs that are detailed by OpenAPI.

EFA offers unique flexibility in supporting multiple IP Fabric topologies based on a BGP underlay with a BGP/EVPN overlay:

- Small Data Center Fabric (non-Clos topology from a single switch pair up to four switch pairs)
- 3-stage Clos (Leaf / Spine)
- 5-stage Clos (Leaf / Spine / Super Spine)

Tenant Network onboarding services are supported on all the topologies, which allows you to create connectivity for devices that are connected to the fabric, such as compute (servers), storage, and any other connectivity needed such as external routers or gateways.

Lifecycle management of the Fabric allows you to add or delete devices to the Fabric after Day 0. Similarly, you can add or remove Tenants as necessary. Key ecosystem integrations streamline Tenant and network provisioning by way of VMware vCenter, Microsoft System Center for Virtual Machine Management (SCVMM) and OpenStack with ML2 and L3 service plugins.

Key Features

The key features of EFA are as follows:

- Conformance to the EVD (Extreme Validated Design for IP Fabric) - <https://www.extremenetworks.com/resources/extreme-validated-design/extreme-ip-fabric-architecture/>
- Ease and speed of provisioning and troubleshooting
- Seamless installation and deployment mechanism
- High performance and low resource utilization with minimal touch points
- Programmable containerized services through an industry-standard Open API (<https://www.openapis.org/>)-based programmable interface
- Easy-to-use CLI commands to manage devices in an IP Fabric and in Tenant Networks
- Ecosystem support for OpenStack, VMware vCenter, and Microsoft Hyper-V

Software Support Matrix

Platforms	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9140	18s.1.01, 18s.1.01a, 18s.1.01c, 18s.1.03	✓				✓
SLX 9240	18s.1.01, 18s.1.01a, 18s.1.01c, 18s.1.03	✓	✓	✓		
SLX 9030	18x.1.00, 18x.1.00a, 18x.1.00b	✓				
SLX 9540	18r.1.00aa, 18r.1.00b, 18r.1.00c, 18r.1.00cc, 20.1.x	✓			✓	
SLX 9640	20.1.x				✓	
SLX 9850	18r.1.00aa, 18r.1.00b, 18r.1.00c		✓	✓		
SLX 9150	20.1.x	✓				✓
SLX 9250	20.1.x	✓	✓			✓

EFA Core Services and Integrations

EFA comprises several core containerized services that interact with each other and with other infrastructure services to provide the core functions of IP Fabric automation.

Service	Description
Asset Service	Provides the secure credential store and deep discovery of physical and logical assets of the managed devices, and publishes the Asset refresh and change events to other services.
Fabric Service	Helps orchestrate and visualize the BGP and EVPN-based 3- and 5-stage CLOS networks.
Tenant Service	Helps manage the Tenants, Tenant Networks, and end points, fully leveraging the knowledge of Assets and the underlying fabric.
Inventory Service	Inventory Service is a REST layer on top of device inventory details, with the ability to filter data based on certain fields. Inventory Service securely stores device credentials in encrypted form and makes them available to different components, such as Fabric Services and Tenant App.

EFA also provides a microservice for each ecosystem integration. This architecture permits rapid development and integration of different ecosystem integrations. Each operates independently to externally integrate while using the same underlying services to interact with the IP Fabric.

Ecosystem	Description
VMware vCenter	The vCenter integration provides connectivity between EFA and vCenter using a REST API as documented in the VI SDK. EFA does not connect to individual ESXi servers. All integration is done through vCenter.
OpenStack	OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter.
Microsoft Hyper-V	The Hyper-V integration supports networking configuration for Hyper-V servers in a datacenter, manual and automated configuration updates when VMs move, and visibility into the VMs and networking resources that are deployed in the Hyper-V setup.

For information about new ecosystem support, see [EFA Ecosystem Integrations](#).

New Features

Fabric Service features

Fabric Service is responsible for automating the N-Stage CLOS underlay and overlay. By default, “overlay” is enabled and can be disabled. Fabric Service exposes the CLI and REST API to clients for automating the N-Stage CLOS underlay and overlay configuration.

New features for this release include:

- Small Data Center Topology (non-Clos) support
- Support for new MCT configuration
- Support for Eco-System Integration; Openstack, VMWare vCenter, Microsoft Hyper-V/SCVMM

Underlay configuration automation includes Interface Configurations (IP Numbered), BGP Underlay for Spine and Leaf, BFD and MCT Configurations.

Overlay configuration automation includes EVPN Configuration and Overlay Gateway Configuration.

Fabric Service is deployed with Inventory Service and Tenant Service.

Tenant Service features

Tenant Service exposes the CLI and REST API to the clients for automating the tenant network configuration on the CLOS and Non-CLOS fabric. Tenant network configuration includes VLAN, BD, VE, EVPN, VTEP, VRF, and Router BGP configuration on the necessary Fabric devices in order to provide Layer 2-extension and Layer 3-extension across the Fabric.

New features for this release include:

- Streamlined Tenant provisioning
- Tenant provisioning on overlay CLOS Fabric
- Tenant provisioning on overlay non-CLOS Fabric

Inventory Service features

Inventory Service is a REST layer on top of device inventory details, with the capability to filter data based on certain fields. Inventory Service will also securely store credentials of devices in encrypted form and make it available to different components like Fabric Services App, and Tenant App.

Inventory service supports device replacement and compare feature

- Device replacement with the same configuration
- Device replacement with different configuration
- Device compare

Inventory Service supports the **execute-cli** option for pushing specific configuration to devices. You can use **execute-cli** to push the same configuration on multiple devices.

EFA Deployment on TPVM

TPVM (Third-Party Virtual Machine) is a general server that resides on Extreme SLX devices.

New for this release is an EFA image that is packaged in the SLX-OS firmware. In this context, EFA leverages the K3S Kubernetes cluster as an underlying infrastructure for the EFA services deployment. You can deploy EFA on an SLX device in the Fabric to manage the Fabric.

Fabric Service Brownfield support

A Brownfield deployment is one in which the installation and configuration of new software must coexist with legacy software systems.

New features for this release include:

- Ability to migrate the devices configured through Embedded Fabric (legacy EFA) to the newer EFA application.
- Ability to migrate the fabric deployed through older version of EFA to newer version of EFA. Here the scenario is user dismantled the older EFA server and installs the newer EFA version or the same EFA version on a different server.

With this feature, you can migrate the fabric being configured fully or partially through the use of the SLX CLI or out-of-band means, provided there are no conflicts with EFA Fabric settings. Brownfield scenarios are not supported for Tenant Services. This is planned for future releases. Fabric Service learns

and fetches the configuration on the devices through the Inventory Service, performs pre-validation checks, and generates appropriate errors for deviations in the configuration.

EFA support for SLX-OS 20.1.x

New features for the release include the following:

- Support for the 20.1.x image running on Leaf, Spine, and Super-Spine devices of CLOS Fabric
- Support for the 20.1.x image running on rack devices of non-CLOS Fabric.
- Support for Fabric and Tenant provisioning on a topology with a mix of pre- 20.1.x images and 20.1.x images.
- Seamless support for upgrade and downgrade of the devices across pre-20.1.x and 20.1.x images to keep the device and application configuration in sync.

For details, see [Software Support Matrix](#).

EFA ecosystem integrations

Administrators can use EFA to integrate with several orchestration ecosystems. EFA provides one-touch integration points with these ecosystems, providing deep insight into the VMs, vSwitches, port groups and hosts, and the translation of these into IP Fabric networking constructs.

New features for the release include the following. For more information, see [EFA Core Services and Integrations](#).

Ecosystem	Integration Highlights								
VMware vCenter	<ul style="list-style-type: none"> • Registration of 1 or more vCenter servers in EFA • Updates for vCenter asset details • Lists of information about vCenter servers • Delete or unregister vCenter servers • Inventory integration • Tenant Service integration • Dynamic updates from vCenter and from EFA services 								
OpenStack	<p>OpenStack plugin package for ML2 and ML3, including :</p> <hr/> <table> <tr> <td>ML2 Plugin (VLAN)</td> <td> <ul style="list-style-type: none"> • CRUD operations on Network and Port • LAG support • Provider Network (default, PT) </td> </tr> <tr> <td>Trunking (VLAN)</td> <td>Trunking using virtio ports</td> </tr> <tr> <td>SRIOV-VF</td> <td>Network Operations using SRIOV-VF Passthrough – Intel/Mellanox</td> </tr> <tr> <td>SRIOV-PF</td> <td>Network Operations using SRIOV-PF Passthrough - Intel/Mellanox</td> </tr> </table>	ML2 Plugin (VLAN)	<ul style="list-style-type: none"> • CRUD operations on Network and Port • LAG support • Provider Network (default, PT) 	Trunking (VLAN)	Trunking using virtio ports	SRIOV-VF	Network Operations using SRIOV-VF Passthrough – Intel/Mellanox	SRIOV-PF	Network Operations using SRIOV-PF Passthrough - Intel/Mellanox
ML2 Plugin (VLAN)	<ul style="list-style-type: none"> • CRUD operations on Network and Port • LAG support • Provider Network (default, PT) 								
Trunking (VLAN)	Trunking using virtio ports								
SRIOV-VF	Network Operations using SRIOV-VF Passthrough – Intel/Mellanox								
SRIOV-PF	Network Operations using SRIOV-PF Passthrough - Intel/Mellanox								

Ecosystem	Integration Highlights	
	Layer 3 (E-W)	East West Traffic using virtio ports (Neutron Router/Router Interface/Subnet CRUD operations)
	VMotion	Virtual Machine Migration.
	BD Support	Support for BD-enabled in Tenant Service
	Multi VIM Support	Multiple Tenant's being managed from OSS
	Multi-Segment Support	Multiple Segments using SRIOV(PF/VF)+ Virtio(DHCP)
	CEP Support	
Microsoft Hyper-V	<ul style="list-style-type: none"> • SCVMM (System Center Virtual Machine Manager) server discovery • SCVMM server update • Periodic polling of registered SCVMM servers • SCVMM server list • SCVMM server delete and de-register • Network event handling 	

SLX firmware download support

With the firmware download with maintenance mode feature, you can download firmware on one or more devices in the IP Fabric with the smallest possible disruption to data path traffic. Both CLOS and Non-CLOS fabrics are supported. Maintenance mode feature is supported only on SLX devices running SLXOS 20.1.x.

New features for this release include:

- Firmware download with maintenance mode
 - Asynchronously launched operation
 - Sanity and pre-install script check
 - Configuration to set convergence timeout, enable, and disable
 - Persist the running configuration so that device configuration and maintenance mode configuration is preserved after reboot
 - Firmware download with **no commit** option to allow restoration of firmware to the previous version
 - During maintenance mode, no configurations will be allowed on the device.
- Firmware host registration, with support for register, update, delete, and list operations
- Firmware download preparation, with support for add, remove, and list operations
- Firmware download execute to start firmware download with maintenance mode asynchronous operation
- Firmware download show, to display a table of devices in the Fabric and their corresponding status

Non-CLOS/Small DC Fabric support

New for this release, support for small DC Fabric offers CLI commands and a REST API that is similar to that of CLOS Fabric.

Non-CLOS topology only has racks interconnected with each other. Non-CLOS topology is used for small scale, VCS-like deployments. The maximum scale is 4 racks and 4*2=8 devices.

This feature is supported with SLX 9140, SLX 9150, and SLX 9250 devices as follows:

- Single rack automation. Each rack consists of two node MCT pair.
- Multi-rack automation
- Multi-homed leaf (MCT)
- BGP neighborship
- Fabric topology view
- Fabric validation and troubleshooting

Prerequisites for SLX Devices

Prerequisite	Description
Supported devices	See Software Support Matrix .
Operating System	See Software Support Matrix .
IP addresses	Management IP addresses configured on all devices
Ports	SLX ports in break-out mode enabled on the devices before EFA can configure the Fabric. For example: <pre>"efa inventory device execute-cli --ip <switch IP> --command "interface ethernet 0/20,shutdown,hardware,connector 0/20, breakout mode 4x10g,interface ethernet 0/20:1-4,no shutdown" -config"</pre>

Prerequisites for the vCenter Controller and Compute Nodes

Minimum System Requirements	Description
vCenter Controller	Windows 2016 with vSphere 6.5.x and 6.7.x
Compute Nodes	ESXi 6.x or above

Prerequisites for the OpenStack Controller and Compute Nodes

Minimum System Requirements	Description
Operating System	Ubuntu 16.04
OpenStack Release	Pike
Extreme Neutron plugin	2.1.0 (on controller) Extreme Neutron plugin is delivered as a Debian package.

Prerequisites for the OpenStack Controller and Compute Nodes

Minimum System Requirements	Description
SCVMM Controller	Windows SCVMM 2016
Compute Nodes	Windows HyperV 2016

OpenStack EFA Neutron Plugin

1. Single-segment Neutron Network creation and deletion:
 - o Internal VLAN network
 - o VLAN provide network (default)
 - o VLAN provider network (default) with external DC edge connectivity
2. Create and delete networks and VMs from Dashboard, Horizon, Neutron CLI, and OpenStack CLI.
3. VM operations such as pause, suspend, shelve, resize, shut, lock, and reboot.

Notes:

- o SLX switch configurations pushed from OpenStack are not automatically saved on the switch. They need to be manually saved on the SLX using the command `copy running-config startup-config`.
- o SLX switch configurations pushed from OpenStack should not be modified either through EFA or direct CLI access to the switch. Doing so will result in the EFA Database going to an inconsistent state. Configurations should only be modified using relevant OpenStack commands.

Fresh Install of EFA on TPVM

Prerequisite: For details about installing and upgrading TPVM, see the *Extreme SLX-OS Software Upgrade Guide, 20.1.1*.

EFA on TPVM is only supported on the SLX 9150 and SLX 9250 platforms.

1. On the SLX where TPVM is planned to be run, verify that TPVM is setup for EFA deployment:
 - o Validate that TPVM is running.
show tpvm status
 - o Validate that TPVM has an assigned IP address.
show tpvm ip-address
 - o Validate that the SSH keys are uploaded.
show tpvm status
 - o Validate that passwordless access is configured.
show tpvm status
2. Enter SLX Linux mode.
start-shell
cd /efaboot
3. Copy the EFA tar file to SLX.
scp <efa-bundle>
4. Deploy EFA on TPVM from SLX shell.
efa deploy

5. Verify status of deployed EFA.

```
# show efa status
```

Fresh Install of EFA on a Server

1. Download the image (*.tar.gz) and untar it.

```
# tar -xzf efa-v2.1.0.tar.gz
```
2. Change directory to a node deployment.

```
# cd efa
```
3. Check for prerequisites.
 - CPU: 4 cores
 - Storage: 50 GB
 - RAM: 8 GB
 - OS: Ubuntu 16.04
4. Run the application installation script.

```
# source deployment.sh
```

Fresh Install of EFA on a VM using OVA

Open Virtual Appliance (OVA) is an OVF file packaged with base image Ubuntu 16.04, storage = 50 G, RAM = 8 GB, and preinstalled with EFA. OVA is also compatible with VMware ESXi servers, so it can be deployed with VMware products too.

Use the OVA image only for new installations. For existing deployments, see [Upgrade EFA on a Server](#) or [Upgrade EFA on TPVM](#).

1. Download the EFA_v2.1.0_<build_number>.ova file.
2. Start Oracle VirtualBox.
3. Choose **File > Import Appliance**.
4. Select the downloaded *.ova file and select **Open**.
5. After importing the OVA file, start the VM.

Notes:

- The user credentials for the OVA installation are:
admin/password: admin/password
root/password: root/dca123
- Log in with “admin” user and then use sudo to run the commands. The new user admin is added in the build.

Upgrade EFA on TPVM

1. Verify that TPVM is set up for EFA deployment:
 - a. Validate that TPVM is running.

```
# show tpvm status
```
 - b. Validate that TPVM has an assigned IP address.

```
# show tpvm ip-address
```
 - c. Validate that the SSH keys are uploaded.

```
# show tpvm status
```

- d. Validate that passwordless access is configured.


```
# show tpvm status
```
2. Determine whether more than one EFA version is available in the SLX `/efaboot` directory.
 - a. If no version is available, the installer stops.
 - b. If more than one version is available, you have the option to pick a version.
 - c. If only one version is available, the installer picks up that version.
3. Determine whether the TPVM already has a version of EFA installed.
 - a. If the same version is already installed, the installer stops.
 - b. If no EFA is installed, the installer continues with installation.
 - c. If a different version is detected, the upgrade or downgrade will continue, depending on the detected version.
4. Copy the EFA tar file to SLX.


```
# start-shell
# cd /efaboot
# scp <efa-bundle>
```
5. Deploy EFA on TPVM from SLX.


```
# efa deploy
```
6. Verify the status of deployed EFA.


```
# show efa status
```

Upgrade EFA on a Server

1. Download the image (*.tar.gz) and untar it.


```
# tar -xzf efa-v2.1.0.tar.gz
```
2. Change directory to a single-node deployment.


```
# cd efa
```
3. Run the deployment script.


```
# source deployment.sh
```

If the previous deployment stack is already running, the script presents the following options:

 - **Remove the current stack:** You can remove the entire stack with this option.
 - **Upgrade or Redeploy:** If you are running the `deployment.sh` script from the new tar-ball, you can upgrade the setup without wiping out the current database volume. Similarly, you have the same option if you rerun the script from the same folder, in which case the stack is only redeployed.
 - **Quit:** No change in the current stack.

Upgrade Notes

- Upon SLX-OS reload or 'tpvm stop'+ 'tpvm start' or Server reload, wait 10 minutes for EFA microservices to be up and running before attempting to invoke EFA commands.
- Use the following k3s commands from TPVM to check if the system is ready. Make sure all PODs are in running state.


```
# k3s kubectl get pods -n efa
```

EFA Database Backup and Restore

Follow these procedures when you need to back up and restore EFA data.

Notes:

- Use sudo privilege to run the backup and restore scripts.
- Use absolute paths. Paths with ~ will not work.
- Before running efa_backup again with the same directory, clean up the previously used directory, because the efa_version file appends the version and then the version comparison fails.

Backup EFA

1. In , create the backup directory.

```
# mkdir /apps/backup
```
2. Run the backup script. Provide the backup directory path when prompted.

```
# bash /apps/efa/efa_backup.sh
```
3. Create a tar file for the backup directory

```
# cd /apps  
# tar cvzf backup.tar.gz backup
```
4. Copy backup data to SLX /efaboot or another server via SCP.

```
# scp admin@<tpvm-ip>:/apps/backup.tar.gz
```

Restore EFA

1. Copy the backup file from the switch /efaboot directory into TPVM.

```
# scp backup.tar.gz admin@<tpvm-ip>:/home/admin/
```
2. Untar the backup data.

```
# tar -xzf backup.tar.gz
```
3. Run the restore script. Provide the untarred data path when prompted.

```
# bash /apps/efa/efa_restore.sh
```

Known Limitations and Issues

1. Need sudo privilege to run efa_backup and efa_restore commands.
2. Backup and restore scripts require absolute paths to be provided and ~ will not work.
3. If same backup directory is reused during efa_backup, backup files will be overwritten.
4. After the Fabric is configured, changes to Fabric settings are not allowed.
5. In a vCenter ecosystem environment, an existing host or DVS (Distributed Virtual Switch) that is present at the time of vCenter registration can be managed by EFA. Any newly added or deleted ESXi host or DVS will not be managed by EFA.
6. In a vCenter ecosystem environment, a DVS Port Group that is present at the time of vCenter registration can be managed by EFA. Any newly added Port Group will not be managed by EFA for CCEP (Cluster Client End Point) port-channels.
7. With Hyper-V ecosystem integration and certain models of older Intel NIC on the compute node, the MAC address is incorrectly discovered and will not work.

Issue #	Issue Summary	Symptom	Workaround
GA-3046	When the switch is not reachable, deleting the network from OpenStack will leave the system inconsistent	EFA and OpenStack databases are inconsistent	Delete the EPG using EFA tenant commands
GA-3304	VM creation in OpenStack is successful when one of the switches in MCT pair is not reachable	VM creation successful but VLAN configuration not applied on switches.	Delete and recreate EPG
GA-3445	"efa openstack debug router delete" fails with the reason "Can not delete the VRF ... as it is configured on the device(s)"	Unable to delete stale router entries that persist in efa	Delete the subnet before deleting router.
GA-3494	PO gets added to EPG on creating VM, when the switch/s are NOT reachable	PO gets added to EPG, but there is no configuration on the device.	Delete and recreate EPG
GA-3219	Missing link information is not updated for some links after ESXi server reboot/inventory update	Same as symptom	Executing EFA inventory update should be avoided when any of the ESXi servers are being rebooted. Wait for all servers to be completely up and operational.
GS-3236	efa inventory list does not show proper information	Efa inventory list does not display ASN/Role/Fabric name for some of the devices	This is a display issue and does not affect functionality. Will be fixed in upcoming release.
GA-3437	vCenter: EPG is not updated with VLAN RANGE	Same as symptom	from vCenter delete the PG and re-create
GA-3445	"efa openstack debug router delete" fails with the reason "Can not delete the VRF .."	Same as symptom	<ol style="list-style-type: none"> 1. Always delete the network from Openstack, avoid deleting on EFA directly 2. use "efa openstack debug network delete " command if there are stale entries
GA-3448	Trunking - Only primary VLAN is deleted from the switch when multiple Trunk VMs are powered-off	Only primary VLAN is deleted from the switch when VMs are powered off	Power-off the Trunk VMs one by one with a one minute time interval between each power-off
GA-3486	Vlans are not deleted when the VM is stopped - instead vlans are getting deleted and created immediately when VM is started	Vlans are not deleted when the VM is stopped	There is no functional impact due to this issue. There is additional delay caused by deletion/creation.
GA-3505	Delete ESXi host event does not delete corresponding EPGs	Same as symptom	Delete EPG directly from EFA in order to cleanup
GA-3506	Datacenter remove event not updating "efa vcenter list"	After datacenter is removed from vCenter, EFA still reports the instance	Delete the stale entry using the command "efa tenant epg delete --name <epg name> --tenant <tenant name>"
GA-3512	EFA errors out with message "Provided VNI: 1006 already"	EPG create fails when delete and add operations are performed in quick succession	After epg delete wait for around 30 seconds of delay before attempting epg create

Issue #	Issue Summary	Symptom	Workaround
	consumed in fabric.” during epg creation		
GA-3540	Fabric provisioning fails due to invalid interface speed	When bringing up large fabrics, configuration fails for some of the nodes	Update inventory and reconfigure fabric.
GA-3595	Creation of a new datacenter on vCenter does not trigger an event in EFA.	Newly created datacenter from vCenter is not reflected in EFA.	Perform vCenter update in order to synchronize EFA with vCenter
GA-3681	One of the EPG mapped to multiple PO is not getting deleted	After vCenter delete, some of the EPGs are still present as stale entries	Use the command “efa tenant epg delete –name <epg name> - - tenant <tenant name>”
GA-3695	Even after switch is removed from the inventory, inventory still lists the switch as present as stale entry	During removal of switch from a fabric, if the switch is unreachable, subsequent inventory deletions will fail for that particular switch	Delete the stale switch entry with the fabric name option. (command: efa inventory device delete –fabric <name>)

EFA Documentation

The *Extreme Fabric Automation Administrator Guide* will be available soon.