



Fabric Engine Release Notes

For Fabric Engine Release 8.9

9037591-00 Rev AB
December 2022



Copyright © 2022 Extreme Networks, Inc.

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

About this Document.....	6
Purpose.....	6
Conventions.....	6
Text Conventions.....	7
Documentation and Training.....	7
Help and Support.....	7
Subscribe to Product Announcements.....	8
Send Feedback.....	8
New in this Release.....	9
New Hardware.....	9
5520 Series Back-to-Front Airflow SKUs.....	9
LRM/MACsec Adapter Support on 5720 Series (LRM use only).....	10
New Transceivers and Components.....	10
New Software Features or Enhancements.....	10
Automatic QoS Priority for MACsec Packets on Intermediate Switches.....	10
Enhanced EDM Help.....	11
Extreme Integrated Application Hosting.....	11
ExtremeCloud IQ Agent Enhancement.....	11
Extreme-Dynamic-ACL Scaling Improvements.....	11
Extreme-Dynamic-Config Improvements.....	12
Factory Default Flag More Granular Options.....	12
Forced Password Change.....	12
Plug and Play Enhancements.....	13
Security Enhancements.....	13
Other Changes.....	14
Boot Loader Image Upgrade on 5520 Series, 5420 Series, and 5320 Series.....	14
CLI changes.....	14
EDM Support for mvpn-isid mac-offset Parameter.....	15
Filenames for this Release.....	15
Upgrade and Downgrade Considerations.....	18
Validated Upgrade Paths.....	18
Validated Upgrade Paths for 5320 Series.....	18
Validated Upgrade Paths for 5420 Series and 5520 Series.....	19
Validated Upgrade Paths for 5720 Series.....	19
Switches That Will Not Use Zero Touch Deployment.....	19
Switches That Will Use Zero Touch Deployment	19
Compatible Fabric IPsec Gateway Versions.....	21
Downgrade Considerations.....	21
Post Upgrade Configuration for Zero Touch Fabric Configuration and Dynamic Nickname Assignment.....	21

Network Requirements.....	22
Zero Touch Fabric Configuration Switch.....	23
Hardware and Software Compatibility.....	25
5320 Series Hardware.....	25
5420 Series Hardware.....	25
5520 Series Hardware.....	26
Operational Notes.....	27
Versatile Interface Module Operational Notes.....	27
Operational Notes for VIM Transceivers.....	28
5720 Series Hardware.....	28
Versatile Interface Module Operational Notes.....	29
Transceivers.....	29
Auto-Negotiation.....	29
Forward Error Correction (FEC).....	29
Power Supply Compatibility.....	30
Scaling.....	31
Layer 2.....	32
Maximum Number of Directed Broadcast Interfaces.....	36
Maximum Number of Microsoft NLB Cluster IP Interfaces.....	36
IP Unicast.....	36
IP Interface Maximums for 5720 Series.....	43
IP Interface Maximums for 5520 Series.....	43
IP Interface Maximums for 5420 Series.....	43
IP Interface Maximums for 5320 Series.....	43
Layer 3 Route Table Size.....	44
Route Scaling.....	44
IP Multicast.....	47
Distributed Virtual Routing (DvR).....	49
Filters, QoS, and Security.....	50
Filter Scaling.....	53
OAM and Diagnostics.....	58
Fabric Scaling.....	61
Maximum Number of Switched UNI Endpoints (C-VID or Untagged Port Bindings).....	65
Maximum Number of SPB Multicast Data I-SIDs.....	67
Number of I-SIDs Supported for the Number of Configured IS-IS Interfaces and Adjacencies (NNIs).....	70
Interoperability Considerations for IS-IS External Metric.....	71
Recommendations.....	72
VRF Scaling.....	72
Important Notices.....	73
ExtremeCloud IQ Support.....	73
Compatibility with ExtremeCloud IQ - Site Engine.....	73
Feature-Based Licensing	74
Memory Usage.....	74
Known Issues and Restrictions.....	75
Known Issues.....	75
Known Issues for 8.9.....	75
Restrictions and Expected Behaviors.....	97

General Restrictions and Expected Behaviors.....	98
Filter Restrictions.....	105
Resolved Issues this Release.....	107
Related Information.....	108
MIB Changes.....	108
Deprecated MIBs.....	108
Modified MIBs.....	108
New MIBs.....	112



About this Document

[Purpose](#) on page 6

[Conventions](#) on page 6

[Documentation and Training](#) on page 7

[Help and Support](#) on page 7

[Send Feedback](#) on page 8

The topics in this section discuss the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document describes important information about this release for platforms that support Extreme Networks Fabric Engine™.

This document includes the following information:

- supported hardware and software
- scaling capabilities
- known issues, including workarounds where appropriate
- known restrictions






Conventions

To help you better understand the information presented in this guide, the following topics describe the formatting conventions used for notes, text, and other elements.

Text Conventions

The following tables list text conventions that can be used throughout this document.

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.

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/.

Help and Support

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

Extreme Portal

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

The Hub

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

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 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.

Send Feedback

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

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

To send feedback, do either of the following:

- 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.



New in this Release

[New Hardware](#) on page 9

[New Software Features or Enhancements](#) on page 10

[Other Changes](#) on page 14

[Filename for this Release](#) on page 15

The following platforms support Fabric Engine 8.9:

- ExtremeSwitching 5320 Series
- ExtremeSwitching 5420 Series
- ExtremeSwitching 5520 Series
- ExtremeSwitching 5720 Series

For MIB-related changes, see [MIB Changes](#) on page 108.



Note

ExtremeSwitching 5420 Series and 5520 Series: Upgrading from an earlier version of VOSS to Fabric Engine 8.6, or later, on these platforms will change the SNMP SysObjectID value. This change might affect SNMP-based management systems. For more information, see this [Knowledge Article](#).

New Hardware

5520 Series Back-to-Front Airflow SKUs

This release supports power supply units (PSU) and fans with back-to-front airflow for the following 5520 Series SKUs:

- 5520-24T-BASE
- 5520-24X-BASE
- 5520-48SE-BASE
- 5520-48T-BASE



Note

These BASE SKUs do not include power supply units or fans; power supply units and fans must be ordered separately.

The following list provides the SKUs for back-to-front airflow components:

- 350W AC PSU — XN-ACPWR-350W-BF

- Fan — 17116

For more information, see [ExtremeSwitching 5520 Series Hardware Installation Guide](#) and [ExtremeSwitching 5520 Series Quick Reference](#).

LRM/MACsec Adapter Support on 5720 Series (LRM use only)

5720-VIM-6YE adds LRM support using an LRM/MACsec adapter, part 10965. The adapter includes two SFP+ host ports and two SFP/SFP+ network/link ports. Host ports use a proprietary 10 Gbps direct attach cable (DAC) to connect to the host (5720-VIM-6YE). You must connect both host ports to trigger PHY initialization and you must insert the DAC in the adapter before you insert the DAC in the VIM.

Both host ports and an auxiliary USB port provide power to the adapter. As a best practice, always use the USB port to provide power.

This release includes new CLI commands to upgrade the internal PHY firmware for the adapter and to view the firmware status.



Note

MACsec functionality of the adapter is not supported, with no impact on MACsec support of the VIM itself.

For more information, see [Two-Port External LRM/MACsec Adapter](#). For the new CLI commands, see [Fabric Engine CLI Commands Reference](#) or [Fabric Engine User Guide](#).

New Transceivers and Components

For optics compatibility, see the [Extreme Optics](#) website.

New Software Features or Enhancements

The following sections describe what is new in this release:


Automatic QoS Priority for MACsec Packets on Intermediate Switches

In certain situations where MACsec encrypted packets traverse intermediate non-MACsec switches, QoS visibility is lost.

This feature uses confidentiality-offset to specify that the first 30 or 50 bytes within the MACsec frame transmit without encryption, thus leaving the 802.1Q VLAN tag p-bits in the clear so that the intermediate switch can differentiate between encrypted traffic. With the 802.1Q p-bits in the clear, internal QoS priority for MACsec packets on intermediate switches can be automatically assigned.

For more information, see [Fabric Engine User Guide](#).

Enhanced EDM Help

Enterprise Device Manager (EDM) has been improved with an online Help feature that supplements the existing Help. A Book icon () in the Navigation pane now provides links to the following items:

- Software Release Notes
- Documentation collections
- Hardware and Software Compatibility Matrices
- Documentation for Extreme optics
- RESTCONF Reference Documentation
- The support portal for Software, MIB, Vulnerability/CVE and Field Notices
- GitHub information for GNS3 images

If you have installed an Extreme-branded transceiver in a port, then you can view information about the transceiver by selecting the Vendor Part Number on the **DDI/SFP** tab of the Port pane.

For more information, see [Fabric Engine User Guide](#).

Extreme Integrated Application Hosting

Beginning with this release, you can use the following existing CLI commands in Global Configuration mode:

- **virtual-service copy-file WORD<1-256> WORD<1-256>**
- **virtual-service WORD<1-128> console**
- **virtual-service WORD <1-128> install package WORD<1-512>**
- **virtual-service WORD<1-128> uninstall**

Procedures are updated to reflect two mode support. Upgrade procedures use Global Configuration mode only to reduce mode changes.

For more information, see [Fabric Engine User Guide](#).

ExtremeCloud IQ Agent Enhancement

The output for the **show application iqagent status** CLI command is updated to provide additional information if IQ Agent is enabled but disconnected. The same information is also available in EDM. This change requires ExtremeCloud IQ Agent 0.5.55 or later.

For more information, see [Fabric Engine CLI Commands Reference](#).

Extreme-Dynamic-ACL Scaling Improvements

The Extreme-Dynamic-ACL RADIUS attribute now supports a list parameter. Use the list parameter to contain ports or masks and group similar ACE commands to avoid the packet limitation. Only one list parameter can exist in one Extreme-Dynamic-ACL. The length of an individual ACE command from an Extreme-Dynamic-ACL VSA message is increased from 128 to 255 characters.

For more information, see [Fabric Engine User Guide](#).

Extreme-Dynamic-Config Improvements

When you use the RADIUS VSA Extreme-Dynamic-Config to activate DHCP Snooping or Dynamic ARP Inspection, this functionality is now only enabled on the RADIUS returned VLAN/I-SID or, in the case of no RADIUS returned VLAN/I-SID, on the untagged VLAN/I-SID already present on the port. In previous releases, the functionality was enabled on all VLANs assigned to the port and this would typically result in warning messages if the onboarding VLAN, a Private VLAN, was also on the Auto-sense port, because DHCP Snooping and Dynamic ARP Inspection are not supported on Private VLANs.

The requirement to have both DHCP Snooping and Dynamic ARP Inspection enabled across all port VLANs is only required if IP Source Guard is to be also enabled on the same port. There is no change from previous releases if you use the RADIUS VSA Extreme-Dynamic-Config to activate IP Source Guard in addition to DHCP Snooping and Dynamic ARP Inspection.

For more information, see [Fabric Engine User Guide](#).

Factory Default Flag More Granular Options

The following list identifies more granular options that enhance the factory default flag behaviors:

- *config-only* — Boots the switch with a blank configuration. This parameter preserves configuration files, primary and secondary configuration file names, user accounts and passwords, digital certificates, IKE/OSPF/IS-IS keys, and SNMP communities. All ports are disabled and assigned to VLAN 1. License files are not removed. Use this parameter as a temporary troubleshooting option to test or investigate if something is wrong with the configuration without permanently removing the configuration files, user accounts, and other preserved items.
- *reset-all-files* — Equivalent to a switch that ships from the factory. The switch has no configuration files, default user accounts, default security mode, Auto-sense-enabled ports, and performs a ZTP+ configuration after reboot. The 30-day factory license is also reset.

If required, you must revoke the license file.



Note

You can also use a new **unconfigure switch** command to achieve the same behavior.

- *zero-touch* — Boots the switch with a default configuration that enables Auto-sense. This parameter resets secure files but keeps the security mode and performs a ZTP+ configuration after reboot. License files are not removed.

For more information, see [Fabric Engine User Guide](#).

Forced Password Change

In this release, the system prompts you to change the admin and read-only user default passwords when you use the **web-server enable** command to enable the web management interface.

For more information, see [Fabric Engine User Guide](#).

Plug and Play Enhancements

This release introduces the following improvements:

- Automatically adjust IS-IS link metrics based on port / LAG speed, which automatically recalculates the Level 1 metric based on the detected link speed.
- Auto-set vim-speed on 25G VIMs based on inserted transceiver type.

The switch now automatically configures the Versatile Interface Module (VIM) speed based on the detected optics, which makes it easier to deploy and to maintain the module.

Auto-set vim-speed is enabled by default.



Note

This feature only applies to 5520 Series and 5720 Series.

- Auto-channelize QSFP+ and QSFP28 ports when QSA adapter or breakout cable is detected and the port operates in Auto-sense mode.

Channelization of these ports occurs automatically when you insert one of the following:

- Quad Small Form-factor Pluggable (QSFP) Plus adapter to Small Form-factor Pluggable (SFP) Plus adapter
- QSFP28 to SFP28 adapter
- QSFP28 to 4xSFP28 passive or active breakout cable
- QSFP+ to 2xSFP+ passive or active breakout cable

This enhancement means that you no longer have to configure channelization on supported ports.

Together these enhancements make it even easier to deploy and use your switch.

For more information, see [Fabric Engine User Guide](#).

Security Enhancements

This release makes the following security-related enhancements:

- Secure syslog automatically reconnects after a connectivity failure

In previous releases, if connectivity failed, the switch disabled the syslog host automatically and you needed to manually retry the connection. Now, if connectivity fails, the syslog host remains enabled and the switch attempts to reconnect with the syslog server every two minutes.

- SSH rekeying applies to Secure Copy (SCP) and Secure File Transfer Protocol (SFTP).

Previously, the SSH rekey data limit and time interval applied only to the SSH server and client.

- If the switch operates in Enhanced Secure Mode (ESM), 3des-cbc and blowfish-cbc encryption types are disabled by default.
- A new log message displays if an SSH packet exceeding 32,768 bytes is received, in both ESM and non ESM modes. In previous releases, the switch silently discarded received SSH packets exceeding 32,768 bytes. For information about log messages, see [Fabric Engine Alarms and Logs Reference](#).

- In both ESM and non ESM modes, the switch limits the supported algorithms for Remote Access Dial-In User Services (RADIUS) Security (RADsec) proxy to the following ciphers:
 - ECDHE-RSA-AES256-GCM-SHA384
 - ECDHE-RSA-AES128-GCM-SHA256
 - ECDHE-RSA-AES256-SHA384
 - ECDHE-RSA-AES128-SHA256
 - DHE-RSA-AES256-SHA256
 - DHE-RSA-AES256-SHA
 - DHE-RSA-AES128-SHA256
 - DHE-RSA-AES128-SHA

**Note**

For RADSec implementations, as a best practice, use radsecproxy version 1.9.1 or later.

- The switch no longer advertises P-192 and P-224 Transport Layer Security (TLS) elliptic curves; the switch advertises P-256, P-384, and P-521. As a best practice with RADsec, manually force the TLS version 1.2 negotiation by adding to the RADsec proxy server configuration file.
- New warning messages display if the switch operates in Enhanced Secure Mode and uses unsecure algorithms. For information about log messages, see [Fabric Engine Alarms and Logs Reference](#).

For more information, see [Fabric Engine User Guide](#).

Other Changes

Boot Loader Image Upgrade on 5520 Series, 5420 Series, and 5320 Series

Fabric Engine 8.9 automatically upgrades the boot loader (uboot) image to the following versions:

- 5520 Series — 2.2.1.14
- 5420 Series — 2.3.2.7
- 5320 Series — 3.4.2.11

When you activate the 8.9 software image, the uboot default information is updated and the switch boots with it. After the reboot, the alternate uboot information is updated.

CLI changes

This release includes the following CLI changes that are not related to new features:

- A **show application openapi** command is available to display the status of openAPI on switches that support ExtremeCloud IQ.
- The output of the **show fulltech** command includes information from the **show khi performance rx-queue** command.
- The output of the **show tech** command is aligned with the output of the **show fulltech** command.
- The **show khi cpp port-statistics** command includes a new parameter, *spbm-internal-ports*, to include internal loopback traffic.

EDM Support for mvpn-isid mac-offset Parameter

This release includes the ability to configure the mac-offset parameter for mvpn-isid from EDM. In previous releases, this functionality was available through CLI only. For more information, see [Fabric Engine User Guide](#).

Filenames for this Release



Important

Do not use Google Chrome or Safari to download software files. Google Chrome can change the file sizes. Safari changes the .tgz extension to .tar.

After you download the software, calculate and verify the md5 checksum. For more information, see [Fabric Engine User Guide](#).

When extracting the software image file, the extraction process appends the software version portion of the extracted filenames to include the final full software version. (For example, extracting **5520.8.2.5.0.voss** results in a software file named **5520.8.2.5.0.GA**.) Ensure that you specify the final full software version (in this case, **8.2.5.0.GA**) when using CLI commands that include the software version, such as activating or removing the software.

The Open Source license text for the switch is included on the product. You can access it by entering the following command in the CLI:

```
more release/w.x.y.z.GA /release/oss-notice.txt
```

where *w.x.y.z* represents a specific release number.

The following tables provide the filenames and sizes for this release.

Table 2: 5320 Series Software Filenames and Sizes

Description	File	Size
Logs reference	5320.8.9.0.0_edoc.tar	64,409,600 bytes
MD5 Checksum files	5320.8.9.0.0.md5	461 bytes
MIB - supported object names	5320.8.9.0.0_mib_sup.txt	1,535,610 bytes
MIB - objects in the OID compile order	5320.8.9.0.0_mib.txt	8,201,788 bytes
MIB - zip file of all MIBs	5320.8.9.0.0_mib.zip	1,225,857 bytes
Open source software - Master copyright file	5320.8.9.0.0_oss-notice.html	2,785,868 bytes
SHA512 Checksum files	5320.8.9.0.0.sha512	1,376 bytes
Software image	5320.8.9.0.0.voss	112,642,924 bytes

Table 2: 5320 Series Software Filenames and Sizes (continued)

Description	File	Size
EDM Help files	FabricEnginev890_HELP_EDM_gzip.zip	4,887,874 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 3: 5420 Series Software Filenames and Sizes

Description	File	Size
Logs reference	5420.8.9.0.0_edoc.tar	64,409,600 bytes
MD5 Checksum files	5420.8.9.0.0.md5	461 bytes
MIB - supported object names	5420.8.9.0.0_mib_sup.txt	1,534,686 bytes
MIB - objects in the OID compile order	5420.8.9.0.0_mib.txt	8,201,788 bytes
MIB - zip file of all MIBs	5420.8.9.0.0_mib.zip	1,225,857 bytes
Open source software - Master copyright file	5420.8.9.0.0_oss-notice.html	2,785,868 bytes
SHA512 Checksum files	5420.8.9.0.0.sha512	1,376 bytes
Software image	5420.8.9.0.0.voss	112,346,412 bytes
EDM Help files	FabricEnginev890_HELP_EDM_gzip.zip	4,887,874 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 4: 5520 Series Software Filenames and Sizes

Description	File	Size
Logs reference	5520.8.9.0.0_edoc.tar	64,409,600 bytes
MD5 Checksum files	5520.8.9.0.0.md5	461 bytes
MIB - supported object names	5520.8.9.0.0_mib_sup.txt	1,535,243 bytes
MIB - objects in the OID compile order	5520.8.9.0.0_mib.txt	8,201,788 bytes
MIB - zip file of all MIBs	5520.8.9.0.0_mib.zip	1,225,857 bytes
Open source software - Master copyright file	5520.8.9.0.0_oss-notice.html	2,785,868 bytes
SHA512 Checksum files	5520.8.9.0.0.sha512	1,376 bytes
Software image	5520.8.9.0.0.voss	117,032,693 bytes

Table 4: 5520 Series Software Filenames and Sizes (continued)

Description	File	Size
EDM Help files	FabricEnginev890_HELP_EDM_gzip.zip	4,887,874 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 5: 5720 Series Software Filenames and Sizes

Description	File	Size
Logs reference	5720.8.9.0.0_edoc.tar	64,409,600 bytes
MD5 Checksum files	5720.8.9.0.0.md5	594 bytes
MIB - supported object names	5720.8.9.0.0_mib_sup.txt	1,540,247 bytes
MIB - objects in the OID compile order	5720.8.9.0.0_mib.txt	8,201,788 bytes
MIB - zip file of all MIBs	5720.8.9.0.0_mib.zip	1,225,857 bytes
Open source software - Master copyright file	5720.8.9.0.0_oss-notice.html	2,785,868 bytes
SHA512 Checksum files	5720.8.9.0.0.sha512	1,701 bytes
Software image	5720.8.9.0.0.voss	318,193,388 bytes
EDM Help files	FabricEnginev890_HELP_EDM_gzip.zip	4,887,874 bytes
Fabric IPsec Gateway	FabricIPSecGW_VM_5.2.0.0.ova	4,034,211,840 bytes
YANG model	restconf_yang.tgz	506,020 bytes
Third Party Virtual Machine (TPVM)	TPVM_Ubuntu20.04_04_14Apr2022.qcow2	4,641,982,464 bytes



Upgrade and Downgrade Considerations

[Validated Upgrade Paths](#) on page 18

[Switches That Will Not Use Zero Touch Deployment](#) on page 19

[Switches That Will Use Zero Touch Deployment](#) on page 19

[Compatible Fabric IPsec Gateway Versions](#) on page 21

[Downgrade Considerations](#) on page 21

[Post Upgrade Configuration for Zero Touch Fabric Configuration and Dynamic Nickname Assignment](#) on page 21

The topics in this section provide information on validated upgrade paths, migration considerations, and compatible software versions.



Note

If a 5420 Series or 5520 Series switch uses DHCP and you did not manually change the host name through the prompt or **sys name** command, applications that are hard-coded with the old host name can be impacted after upgrade from a VOSS release to Fabric Engine 8.6 or later. As a workaround, change the system name or prompt back to `voss<mac-address>`.

See the [Fabric Engine User Guide](#) for detailed image management procedures that includes information about the following specific upgrade considerations:

- IPv6:
 - Notes for systems using IPv6 static neighbors

Upgrade switches using one of the options in the following sections:

- [Switches That Will Not Use Zero Touch Deployment](#) on page 19
- [Switches That Will Use Zero Touch Deployment](#) on page 19

Validated Upgrade Paths

This section identifies the software releases for which upgrades to this release have been validated.

Validated Upgrade Paths for 5320 Series

- Fabric Engine 8.8.x to Fabric Engine 8.9
- Fabric Engine 8.7.x to Fabric Engine 8.9
- Fabric Engine 8.6.x to Fabric Engine 8.9

Validated Upgrade Paths for 5420 Series and 5520 Series

- Fabric Engine 8.8.x to Fabric Engine 8.9
- Fabric Engine 8.7.x to Fabric Engine 8.9
- Fabric Engine 8.6.x to Fabric Engine 8.9
- VOSS 8.5.x to Fabric Engine 8.9



Note

For any pre-8.5.0.0 versions, an intermediate upgrade is required.

Validated Upgrade Paths for 5720 Series

- Fabric Engine 8.8.x to Fabric Engine 8.9
- Fabric Engine 8.7.x to Fabric Engine 8.9

Switches That Will Not Use Zero Touch Deployment

Switches that will not use Zero Touch Deployment with ExtremeCloud™ IQ or ZTP+ with ExtremeCloud IQ - Site Engine should upgrade to this release by performing these steps:

1. Upgrade to this release from one of the previously described releases, see [Validated Upgrade Paths](#) on page 18.
2. Continue to use the previous switch configuration.

Switches That Will Use Zero Touch Deployment

Switches that will use Zero Touch Deployment with ExtremeCloud IQ or ZTP+ with ExtremeCloud IQ - Site Engine should upgrade to this release by performing the following steps:



Important

When you perform these steps, any prior configuration for this switch is lost.

You do not need to complete this procedure for switches that are already managed by ExtremeCloud IQ or ExtremeCloud IQ - Site Engine; use the upgrade functionality available in ExtremeCloud IQ or ExtremeCloud IQ - Site Engine.

1. Upgrade to this release from one of the previously described releases, see [Validated Upgrade Paths](#) on page 18.
2. Ensure the switch boots without a configuration file. To ensure the switch boots without a configuration file, perform one of the following actions:
 - Rename existing primary and secondary configuration files. Use the **mv** command to rename the existing configuration files. For example, **mv config.cfg config.cfg.backup**.

This is the preferred option because it ensures that the primary and secondary files are removed while making a backup of them at the same time. This option also ensures that the switch uses the default config.cfg file for the final configuration after it has successfully onboarded.

- Delete the existing primary and secondary configuration files. Create a backup of these files before you delete them.

- Boot from non-existent configuration files. Use the **boot config choice** command to configure the primary and backup configuration files to reference files that do not exist on the switch:

```
boot config choice primary config-file nonexistent1.cfg
```

```
boot config choice primary backup-config-file nonexistent2.cfg
```

This option also works, however, after the switch has successfully onboarded, it does not use the default config.cfg file but uses the alternative configuration file name provided instead, which might not be desired.

3. Reboot the switch.

Performing these steps results in a switch with a Zero Touch Deployment configuration with the following characteristics:

- The ssh and sshd boot configuration flags are enabled by default.
- All ports are Private VLAN isolated ports.
- VLAN 4048 is created as an *onboarding-vlan* for host-only connectivity for In Band management. All front panel ports are members of VLAN 4048.
- In Band management is enabled.
- Dynamic Host Configuration Protocol (DHCP) client requests are cycled between In Band and Out of Band ports, except on the 5320 Series. 5320 Series supports In Band management only.
- If the switch resets after the IP address is obtained from the DHCP Server, the entire DHCP process does not need to be repeated. Instead, the switch can directly send the DHCP Request to the DHCP Server for the IP stored in the /intflash/dhcp/dhclient.leases file.
- Out of Band management is enabled, except on the 5320 Series. 5320 Series supports In Band management only.
- All ports are administratively enabled.
- IQAgent is enabled by default.
- Zero Touch Provisioning Plus (ZTP+) for ExtremeCloud IQ - Site Engine onboarding is enabled by default.
- Zero Touch Fabric Configuration is initiated.
- After the Zero Touch Fabric establishes successfully, the onboarding VLAN 4048 is automatically assigned to onboarding I-SID 15999999.

After the switch reboots in the Zero Touch Deployment configuration, the DHCP client and ExtremeCloud IQ Agent are enabled. The DHCP client obtains an IP address for the switch, DNS discovery is used to discover a Domain Name Server, and the switch attempts to connect to ExtremeCloud IQ and ExtremeCloud IQ - Site Engine.

All switches also receive a Zero Touch Fabric Configuration. For more information, see [Fabric Engine User Guide](#).

Compatible Fabric IPsec Gateway Versions



Note

This section only applies to 5720-24MXW and 5720-48MXW. For more information about feature support, see [Fabric Engine and VOSS Feature Support Matrix](#).

The OVA image for the Fabric IPsec Gateway is posted with the image file for each network operating system (NOS) release.

For more information about image files in this release, see [Filenames for this Release](#) on page 15. For virtual service upgrade instructions, see [Fabric Engine User Guide](#).

Only use the Fabric IPsec Gateway image version that is posted with the NOS release image.



Note

Upgrade the switch software image before you upgrade the Fabric IPsec Gateway image.

Downgrade Considerations

Save a backup copy of your switch configuration before upgrading to new release. New releases contain significant enhancements, which cannot be used in previous software versions. Downgrading to an earlier release will require a compatible configuration file.

For devices running VOSS 8.3, Fabric Engine 8.6, or later, that connect to ExtremeCloud IQ using ExtremeCloud IQ Agent versions 0.4.0 or higher, you cannot downgrade to VOSS 8.2.x and connect to the cloud automatically. After you downgrade to VOSS 8.2.x, you lose connectivity to ExtremeCloud IQ so you must install a VOSS 8.2.x compatible ExtremeCloud IQ Agent version to re-establish connectivity.

Contact support for assistance with installation of the VOSS 8.2.x compatible ExtremeCloud IQ Agent version. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.



Note

Prior to Fabric Engine 8.6, 5520 Series and 5420 Series platforms ran VOSS. VOSS support ends for these platforms with VOSS 8.5.x.

For information about how to reinstall ExtremeCloud IQ Agent firmware, see [Fabric Engine User Guide](#).

Post Upgrade Configuration for Zero Touch Fabric Configuration and Dynamic Nickname Assignment



Note

In this section, a Zero Touch Fabric release refers to any of the following: VOSS 8.3, Fabric Engine 8.6, or later releases.

The switch initiates Zero Touch Fabric Configuration if you boot without a configuration file.

To add new Zero Touch Fabric Configuration devices or implement Zero Touch Fabric Configuration on existing devices, the network requires a nickname server and reachability to the DHCP server

and, optionally, ExtremeCloud IQ servers or ExtremeCloud IQ - Site Engine. How you implement this depends on if the network is a new deployment, or an existing Fabric network that you upgrade. In a new deployment, you can meet the network requirements with one node, known as a seed node. In an existing network, functions can already exist on different nodes.

For more details on Zero Touch Fabric Configuration, see [Fabric Engine User Guide](#).



Important

Not all upgrade paths are validated by Extreme Networks for each new software release. To understand the validated upgrade paths, see [Validated Upgrade Paths](#) on page 18.

Network Requirements

The following list identifies the network requirements before you add new Zero Touch Fabric Configuration devices or implement Zero Touch Fabric Configuration on existing devices:

- You must configure a node as the nickname server, if one does not already exist. This node can be anywhere in the SPB Fabric IS-IS area.
- The DHCP server must be reachable by the remote nodes:
 - In an existing network, the DHCP server can be anywhere in the network. If the DHCP server is on a different IP subnet from the onboarding I-SID, configure DHCP Relay functionality on the existing IP interface of VLAN 4048 with I-SID 15999999.
 - If the DHCP server is on the same subnet as the onboarding I-SID, configure the port facing the DHCP server as private-vlan promiscuous, using Private VLAN 4048, if the new DHCP snooping port feature does not have the promiscuous port configured automatically. This VLAN and the Auto-sense onboarding I-SID are created automatically on a newly deployed device.
- In this release, ports send Fabric Connect LLDP TLVs regardless of the Auto-sense configuration, which means these devices can establish adjacencies with other devices that run a Zero Touch Fabric release, and use either Auto-sense or static NNI configuration.

In an existing network that includes devices that run a version of VOSS earlier than 8.3, you must manually configure the NNI. Because the port running in the earlier release does not send Fabric Connect LLDP TLVs, an adjacency with a Zero Touch Fabric release node does not form automatically.

For Zero Touch Fabric Configuration to work when a new switch that runs a Zero Touch Fabric release, connects to a switch on an existing Fabric, upgrade at least the existing Fabric switches to a Zero Touch Fabric release first.

- Some SPB deployments use Ethertype 0x88a8 but many use 0x8100. Zero Touch Fabric Configuration works with existing networks that use either value as long as the existing switches that connect to the new switches run a Zero Touch Fabric release.

Zero Touch Fabric Configuration Switch



Important

If you deploy a Fabric-capable switch with Auto-sense enabled, the switch interacts with existing switches that support Fabric Attach (FA). If an existing FA Proxy switch does not have FA server connectivity established yet, it will form an FA connectivity to the newly connected VOSS (8.3 or later) or Fabric Engine, switch as it announces itself as an FA server. To avoid unintended FA connectivity, disable Auto-sense using the **no auto-sense enable** command on the relevant ports.

On switches (upgraded existing or newly deployed) where you want to initiate Zero Touch Fabric Configuration, perform the following tasks:

1. Upgrade to a Zero Touch Fabric release, if the device is not a new deployment already running a Zero Touch Fabric release. As a best practice, upgrade to a Fabric Engine release. For a new deployment of universal hardware, ensure the network operating system (NOS) is Fabric Engine.
2. On upgraded existing switches, ensure the switch boots without a configuration file. The switch joins the network as an end host. To ensure the switch boots without a configuration file, perform one of the following actions:

- Rename existing primary and secondary configuration files. Use the **mv** command to rename the existing configuration files. For example, **mv config.cfg config.cfg.backup**.

This is the preferred option because it ensures that the primary and secondary files are removed while making a backup of them at the same time. This option also ensures that the switch uses the default config.cfg file for the final configuration after it has successfully onboarded.

- Delete the existing primary and secondary configuration files. Create a backup of these files before you delete them.
- Boot from non-existent configuration files. Use the **boot config choice** command to configure the primary and backup configuration files to reference files that do not exist on the switch:

```
boot config choice primary config-file nonexistent1.cfg
```

```
boot config choice primary backup-config-file nonexistent2.cfg
```

This option also works, however, after the switch has successfully onboarded, it does not use the default config.cfg file but uses the alternative configuration file name provided instead, which might not be desired.

3. The switch creates a Zero Touch Deployment configuration to onboard the switch, including the following Zero Touch Fabric Configuration items:



Note

For more details on Zero Touch Deployment, see [Fabric Engine User Guide](#).

- Creates private VLAN 4048.
- Enables SPBM.
- Creates SPBM instance 1.
- Creates default backbone VLANs (B-VLAN) (4051 and 4052).

- Creates manual area 00.1515.fee1.900d.1515.fee1.900d.

**Note**

The B-VLAN and manual area configuration values are not compulsory. This remote switch can attach to a Fabric core that does not match these values because the Auto-sense functionality dynamically learns the B-VLANs and manual area in use in the Fabric core from the connected seed node using LLDP.

- Creates the onboarding I-SID 15999999.
- Assigns the onboarding I-SID to private VLAN 4048 and also includes the management VLAN.

**Note**

As a best practice, use the onboarding I-SID for onboarding purposes and, whenever possible, configure a management VLAN or management CLIP on a different I-SID after the onboarding procedures have been successfully completed.

- Enables Auto-sense on all ports.
 - Configures Auto-sense access ports and Layer 2 trusted Auto-sense ports.
 - Enables IS-IS globally.
 - With Auto-sense, ports on a switch can detect whether they connect to an SPB device, a Fabric Attach (FA) client, FA Proxy, Voice IP devices, or an undefined host, and then make the necessary configuration.
4. If the seed node uses Auto-sense IS-IS Authentication, configure the remote switch to use the same authentication type and key as the seed node.
 5. The switch joins the Fabric.
 6. The nickname server dynamically assigns an SPBM nickname.
 7. After the Zero Touch Fabric establishes successfully, the switch attempts to acquire an IP address on the onboarding VLAN and I-SID using DHCP. When the DHCP client obtains an IP address for the switch, the switch automatically attempts to connect to ExtremeCloud IQ and Extreme Management Center or ExtremeCloud IQ - Site Engine.



Hardware and Software Compatibility

- [5320 Series Hardware](#) on page 25
- [5420 Series Hardware](#) on page 25
- [5520 Series Hardware](#) on page 26
- [5720 Series Hardware](#) on page 28
- [Transceivers](#) on page 29
- [Power Supply Compatibility](#) on page 30

The topics in this section list the software compatibility for hardware platforms.

5320 Series Hardware

5320 Series is a universal hardware product that supports more than one Network Operating System (NOS) personality. For information about NOS personalities, see [Fabric Engine User Guide](#).

Table 6: Switch models

Model	Initial release	Supported new Fabric Engine feature release			
		8.6	8.7	8.8	8.9
5320-16P-4XE	Fabric Engine 8.6.1	N	Y	Y	Y
5320-16P-4XE-DC	Fabric Engine 8.6.1	N	Y	Y	Y
5320-24P-8XE	Fabric Engine 8.6	Y	Y	Y	Y
5320-24T-8XE	Fabric Engine 8.6	Y	Y	Y	Y
5320-48P-8XE	Fabric Engine 8.6	Y	Y	Y	Y
5320-48T-8XE	Fabric Engine 8.6	Y	Y	Y	Y

5420 Series Hardware

5420 Series is a universal hardware product that supports more than one Network Operating System (NOS) personality. For information about NOS personalities, see [Fabric Engine User Guide](#).

**Note**

Prior to Fabric Engine 8.6, 5420 Series platforms ran VOSS. VOSS support ends for these platforms with VOSS 8.5.x.

Table 7: Switch models

Model	Initial release	Supported new Fabric Engine feature release				
		VOSS 8.5	8.6	8.7	8.8	8.9
5420F-24T-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-8W-16P-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-24P-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-48T-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-16MW-32P-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-24S-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-16W-32P-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-48P-4XE	VOSS 8.4	Y	Y	Y	Y	Y
5420F-48P-4XL	VOSS 8.4	Y	Y	Y	Y	Y
5420M-24T-4YE	VOSS 8.4	Y	Y	Y	Y	Y
5420M-24W-4YE	VOSS 8.4	Y	Y	Y	Y	Y
5420M-48T-4YE	VOSS 8.4	Y	Y	Y	Y	Y
5420M-48W-4YE	VOSS 8.4	Y	Y	Y	Y	Y
5420M-16MW-32P-4YE	VOSS 8.4	Y	Y	Y	Y	Y

5520 Series Hardware

5520 Series is a universal hardware product that supports more than one Network Operating System (NOS) personality. For information about NOS personalities, see [Fabric Engine User Guide](#).

**Note**

Prior to Fabric Engine 8.6, 5520 Series platforms ran VOSS. VOSS support ends for these platforms with VOSS 8.5.x.

Table 8: Switch models

Model	Initial release	Supported new Fabric Engine feature release				
		VOSS 8.5	8.6	8.7	8.8	8.9
5520-24T	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-24W	VOSS 8.2.5	Y	Y	Y	Y	Y

Table 8: Switch models (continued)

Model	Initial release	Supported new Fabric Engine feature release				
		VOSS 8.5	8.6	8.7	8.8	8.9
5520-48T	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-48W	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-12MW-36W	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-24X	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-48SE	VOSS 8.2.5	Y	Y	Y	Y	Y

**Note**

Ensure the switch runs, at a minimum, the noted initial software release before you install a VIM.

Table 9: Versatile Interface Modules (VIMs)

Model	Initial release	Supported new Fabric Engine feature release				
		VOSS 8.5	8.6	8.7	8.8	8.9
5520-VIM-4X	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-VIM-4XE	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-VIM-4YE	VOSS 8.2.5	Y	Y	Y	Y	Y

Operational Notes

The 5520-24T, 5520-24X, 5520-48SE, and 5520-48T models require a minimum of Fabric Engine 8.9 to support power supply units and fans with back-to-front airflow.

Versatile Interface Module Operational Notes

The following table summarizes the operational capabilities of the various VIMs:

Table 10: 5520-VIM Matrix

	5520-VIM-4X	5520-VIM-4XE	5520-VIM-4YE
Operational speeds	1Gbps & 10Gbps	1Gbps & 10Gbps	10Gbps & 25Gbps
PHY present	No	Yes	Yes
1000BASE-T & 10GBASE-T	10GBASE-T only	Both	10GBASE-T only
Mixed speeds	1Gbps & 10Gbps	1Gbps & 10Gbps	Mixed speeds not supported
1G Auto-negotiation	Disabled	Disabled	Disabled
10G Auto-negotiation	Disabled	Disabled	Disabled
25G Auto-negotiation			Enabled for DAC Disabled for Fiber

Table 10: 5520-VIM Matrix (continued)

	5520-VIM-4X	5520-VIM-4XE	5520-VIM-4YE
FEC	Not supported	Not supported	Auto-FEC enabled for DAC and Fiber
MACsec	Not supported	128/256 bit	128/256 bit

Operational Notes for VIM Transceivers

The IEEE 802.3by requirement for 25 Gb is that any transceiver or DAC 3 meters or longer, requires the use of forward error correction (FEC).

If you use an unsupported 25 Gb transceiver, you might experience CRC or link flap errors.

5720 Series Hardware

5720 Series is a universal hardware product that supports more than one Network Operating System (NOS) personality. For information about NOS personalities, see [Fabric Engine User Guide](#).

Table 11: Switch models

Model	Initial release	Supported new Fabric Engine feature release	
		8.8	8.9
5720-24MW	Fabric Engine 8.7	Y	Y
5720-24MXW	Fabric Engine 8.7	Y	Y
5720-48MW	Fabric Engine 8.7	Y	Y
5720-48MXW	Fabric Engine 8.7	Y	Y

**Note**

Ensure the switch runs, at a minimum, the noted initial software release before you install a VIM.

Table 12: Versatile Interface Modules (VIMs)

Model	Initial release	Supported new Fabric Engine feature release	
		8.8	8.9
5720-VIM-2CE	Fabric Engine 8.7	Y	Y
5720-VIM-6YE	Fabric Engine 8.7	Y	Y

Versatile Interface Module Operational Notes

The following table summarizes the operational capabilities of the various VIMs:

Table 13: 5720-VIM Matrix

	5720-VIM-2CE	5720-VIM-6YE
Operational speeds	10/25/40/100Gbps	1/10/25Gbps
PHY present	Yes	Yes
1000BASE-T & 10GBASE-T	10GBASE-T only	Both
Mixed speeds	10/25/40Gbps	1/10/25Gbps
1G Auto-negotiation	Not supported	Not supported
10G Auto-negotiation	Not supported	Not supported
25G Auto-negotiation	Supported	Supported
FEC	Supports CL74/CL91	Supports CL74/CL91
MACsec	128/256 bit	128/256 bit

Transceivers

The software allows the use of transceivers and direct attach cables from any vendor, which means that the switch will bring up the port operationally when using any transceiver. Extreme Networks does not provide support for operational issues related to the use of non-Extreme Networks branded transceivers and direct attached cables used in the switches.

To find product descriptions and compatibility information for optical transceivers and components, visit the [Extreme Optics](#) website.

Auto-Negotiation

Use auto-negotiation to enable the device to automatically negotiate the best common data rate and duplex mode to use between two auto-negotiation-capable Ethernet devices.

When you use a 1 Gb SFP transceiver on a 10 Gb SFP+ port, ensure that auto-negotiation is enabled.

For 1000BASE-T SFP transceivers, the best practice is to perform custom auto-negotiation at the remote native copper port. This can prevent connections from failing if the speed or duplex negotiation changes.

Forward Error Correction (FEC)

Forward Error Correction (FEC) is a method of obtaining error control in data transmission over an unreliable or noisy channel in which the source (transmitter) encodes the data in a redundant way by using an error correcting code (ECC). This redundancy enables a destination (receiver) to detect a limited number of errors and correct them without requiring a re-transmission.

For more information about FEC, see [Fabric Engine User Guide](#).

Power Supply Compatibility

You can use certain power supplies in more than one platform.

For more specific information on each power supply, see the following documents:

- [ExtremeSwitching 5320 Series Hardware Installation Guide](#)
- [ExtremeSwitching 5420 Series Hardware Installation Guide](#)
- [ExtremeSwitching 5520 Series Hardware Installation Guide](#)
- [ExtremeSwitching 5720 Series Hardware Installation Guide](#)



Scaling

[Layer 2 on page 32](#)

[IP Unicast on page 36](#)

[Layer 3 Route Table Size on page 44](#)

[IP Multicast on page 47](#)

[Distributed Virtual Routing \(DvR\) on page 49](#)

[Filters, QoS, and Security on page 50](#)

[OAM and Diagnostics on page 58](#)

[Fabric Scaling on page 61](#)

[VRF Scaling on page 72](#)

This section documents scaling capabilities of the universal hardware platforms.

The scaling and performance information shown in the following tables is provided for the purpose of assisting with network design. It is recommended that network architects and administrators design and manage networks with an appropriate level of network scaling “head room.” The scaling and performance figures provided have been verified using specific network topologies using limited switch configurations. There is no guarantee that the scaling and performance figures shown are applicable to all network topologies and switch configurations and are provided as a realistic estimation only. If you experience scaling and performance characteristics that you feel are sufficiently below what has been documented, contact Extreme Networks technical support for additional assistance.



Note

If your switch uses Advanced Feature Bandwidth Reservation in Full Feature mode, this affects scaling information that is based on the number of available ports. If you enable the boot configuration flag for this feature, remember to deduct the number of reserved ports from the documented scaling maximum. Not all hardware platforms require this feature to provide full feature support. For more information, see [Fabric Engine User Guide](#).

Layer 2

Table 14: Layer 2 Maximums

Attribute	Product	Maximum number supported
MAC table size (without SPBM)	5320 Series	32,000
	5420 Series	32,000 for 5420F Series models 64,000 for 5420M Series models
	5520 Series	81,920
	5720 Series	164,000 for 5720MXW models 100,000 for 5720MW models
MAC table size (with SPBM)	5320 Series	16,000
	5420 Series	16,000 for 5420F Series models 32,000 for 5420M Series models
	5520 Series	40,960
	5720 Series	82,000 for 5720MXW models 50,000 for 5720MW models
Endpoint Tracking MAC addresses per switch	5320 Series	n/a
	5420 Series	n/a
	5520 Series	8,000
	5720 Series	8,000
Directed Broadcast interfaces	5320 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 36.
	5420 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 36.
	5520 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 36.
	5720 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 36.

Table 14: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
Port-based VLANs Note: When you use Flex-UNI functionality, you can use the complete range from 1 to 4096 for port VLAN IDs.	5320 Series	4,059
	5420 Series	4,059
	5520 Series	4,059
	5720 Series	4,059
Private VLANs	5320 Series	See Table 15 on page 35
	5420 Series	See Table 15 on page 35
	5520 Series	See Table 15 on page 35
	5720 Series	See Table 15 on page 35
Protocol-based VLANs (IPv6 only)	5320 Series	1
	5420 Series	1
	5520 Series	1
	5720 Series	1
RSTP instances	5320 Series	1
	5420 Series	1
	5520 Series	1
	5720 Series	1
MSTP instances	5320 Series	12
	5420 Series	12
	5520 Series	12
	5720 Series	12
LACP aggregators	5320 Series	48-port models: 56 24-port models: 32 16-port models: 20
	5420 Series	56 (48 fixed ports, 4 Universal Ethernet ports, 4 SFP-DD ports)
	5520 Series	48-port models: up to 60 with channelization 24-port models: up to 36 with channelization
	5720 Series	64

Table 14: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
Ports per LACP aggregator	5320 Series	8 active
	5420 Series	8 active
	5520 Series	8 active
	5720 Series	8 active
MLT groups	5320 Series	48-port models: 56 24-port models: 32 16-port models: 20
	5420 Series	56 (48 fixed ports, 4 Universal Ethernet ports, 4 SFP-DD ports)
	5520 Series	48-port models: up to 60 with channelization 24-port models: up to 36 with channelization
	5720 Series	64
Ports per MLT group	5320 Series	8 active
	5420 Series	8 active
	5520 Series	8
	5720 Series	8
Link State Tracking (LST) groups	5320 Series	48
	5420 Series	48
	5520 Series	48
	5720 Series	48
Interfaces per LST group	5320 Series	48-port models: 9 upstream/128 downstream 16- and 24-port models: 8 upstream/128 downstream
	5420 Series	8 upstream 128 downstream
	5520 Series	8 upstream 128 downstream
	5720 Series	8 upstream 128 downstream
SLPP VLANs	5320 Series	128
	5420 Series	128
	5520 Series	128
	5720 Series	500

Table 14: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
VLACP interfaces	5320 Series	48-port models: 56 24-port models: 32 16-port models: 20
	5420 Series	56 (48 fixed ports, 4 Universal Ethernet ports, 4 SFP-DD ports)
	5520 Series	48-port models: up to 60 with channelization 24-port models: up to 36 with channelization
	5720 Series	64 with no SPB mode: up to 56 with SPBM mode with the channelization enabled when using 5720-VIM-2CE. 64 with no VIM: up to 54 with 5720-VIM-6YE.
Microsoft NLB cluster IP interfaces	5320 Series	Not supported
	5420 Series	Not supported
	5520 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 36.
	5720 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 36.

The number of Private VLANs/Layer 2 E-Tree varies depending on the number of private VLAN trunk ports as members. The following table provides the maximum numbers.

Table 15: Private VLAN and Layer 2 E-Tree maximums

Platform	Total Private VLANs and Layer 2 E-Tree with 2 Private VLAN trunk ports	Total Private VLANs and Layer 2 E-Tree with 4 Private VLAN trunk ports
5320 16- and 24-port models	40	20
5320 48-port models	100	50
5420 Series	100	50
5520 Series	200	100
5720 Series	200	100

Maximum Number of Directed Broadcast Interfaces

The number of Directed Broadcast interfaces must be less than or equal to 200. However, if you configure VLANs with both NLB and Directed Broadcast, you can only scale up to 100 VLANs.

Maximum Number of Microsoft NLB Cluster IP Interfaces

The number of NLB cluster IP interfaces multiplied by the number of configured clusters must be less than or equal to 200. The number of NLB cluster IP interfaces is the key, not the number of VLANs. You can configure 1 VLAN with up to 200 NLB cluster IP interfaces or configure up to 200 VLANs with 1 NLB cluster IP interface per VLAN.

For example: 1 virtual interface per cluster x 200 clusters = 200 or 2 virtual interfaces per cluster x 100 clusters = 200

However, if you configure VLANs with both NLB and Directed Broadcast, you can only scale up to 100 VLANs assuming there is only 1 NLB cluster IP interface per VLAN.

IP Unicast

Table 16: IP Unicast Maximums

Attribute	Product	Maximum number supported
IP interfaces (IPv4 or IPv6 or IPv4+IPv6)	5320 Series	248 See IP Interface Maximums for 5320 Series on page 43.
	5420 Series	248 See IP Interface Maximums for 5420 Series on page 43.
	5520 Series	500 See IP Interface Maximums for 5520 Series on page 43.
	5720 Series	1000 See IP Interface Maximums for 5720 Series on page 43.

Table 16: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
VRRP interfaces (IPv4 or IPv6)	5320 Series	48-port models: 124 16- and 24-port models: 64 See IP Interface Maximums for 5320 Series on page 43.
	5420 Series	124 See IP Interface Maximums for 5420 Series on page 43.
	5520 Series	252 See IP Interface Maximums for 5520 Series on page 43.
	5720 Series	500 See IP Interface Maximums for 5720 Series on page 43.
Routed Split Multi-Link Trunking (RSMLT) interfaces (IPv4 or IPv6 or IPv4+IPv6)	5320 Series	n/a
	5420 Series	124 See IP Interface Maximums for 5420 Series on page 43.
	5520 Series	499 See IP Interface Maximums for 5520 Series on page 43.
	5720 Series	500 See IP Interface Maximums for 5720 Series on page 43.
VRRP interfaces with fast timers (200ms) - IPv4/IPv6	5320 Series	24
	5420 Series	24
	5520 Series	24
	5720 Series	24 See IP Interface Maximums for 5720 Series on page 43.
ECMP groups/paths per group	5320 Series	48-port models: 64/8 16- and 24-port models:32/8
	5420 Series	64/8
	5520 Series	256/8
	5720 Series	2,048/8

Table 16: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
OSPF v2/v3 interfaces	5320 Series	48-port models: 50 16- and 24-port models: 1
	5420 Series	50
	5520 Series	100
	5720 Series	65
OSPF v2/v3 neighbors (adjacencies)	5320 Series	50
	5420 Series	50
	5520 Series	100
	5720 Series	500
OSPF areas	5320 Series	48-port models: 12 16- and 24-port models: 4
	5420 Series	12 for the switch
	5520 Series	12 for each VRF 80 for the switch
	5720 Series	12 for each VRF 80 for the switch
IPv4 ARP table	5320 Series	48-port models: 15,000 16- and 24-port models: 8,000
	5420 Series	15,000 for 5420F Series models 24,000 for 5420M Series models
	5520 Series	16,000
	5720 Series	24,000 for 5720 MW Series models 64,000 for 5720 MXW Series models
IPv4 CLIP interfaces	5320 Series	64
	5420 Series	64
	5520 Series	64
	5720 Series	64
IPv4 RIP interfaces	5320 Series	50
	5420 Series	50
	5520 Series	100
	5720 Series	200

Table 16: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
IPv4 BGP peers	5320 Series	8
	5420 Series	8
	5520 Series	16
	5720 Series	256
IPv4 VRFs with iBGP	5320 Series	48-port models: 8 16- and 24-port models: 1
	5420 Series	8
	5520 Series	16
	5720 Series	16
IPv4/IPv6 VRF instances For additional information, see VRF Scaling on page 72.	5320 Series	48-port models: 64 16- and 24-port models: 1 See IP Interface Maximums for 5320 Series on page 43.
	5420 Series	64 See IP Interface Maximums for 5420 Series on page 43.
	5520 Series	256 including mgmt VRF and GRT See IP Interface Maximums for 5520 Series on page 43.
	5720 Series	256 See IP Interface Maximums for 5720 Series on page 43.
IPv4 static ARP entries	5320 Series	48-port models: 1,000 per VRF/5,000 per switch 16- and 24-port models: 1,000 per switch
	5420 Series	1,000 per VRF 5,000 per switch
	5520 Series	2,000 for each VRF 10,000 for the switch
	5720 Series	2,000 for each VRF 10,000 for the switch

Table 16: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
IPv4 static routes	5320 Series	48-port models: 500 per VRF/2,500 per switch 16- and 24-port models: 500 per switch
	5420 Series	500 per VRF 2500 per switch
	5520 Series	1,000 for each VRF 5,000 for the switch
	5720 Series	1,000 for each VRF 5,000 for the switch
IPv4 route policies	5320 Series	48-port models: 50 per VRF/500 per switch 16- and 24-port models: 500 per switch
	5420 Series	50 per VRF 500 per switch
	5520 Series	500 for each VRF 5,000 for the switch
	5720 Series	500 for each VRF 5,000 for the switch
IPv4 UDP forwarding entries	5320 Series	128
	5420 Series	128
	5520 Series	256
	5720 Series	512
IPv4 DHCP Relay forwarding entries	5320 Series	248
	5420 Series	248
	5520 Series	512
	5720 Series	2048
IPv6 DHCP Snoop entries in Source Binding Table	5320 Series	48-port models: 513 16- and 24-port models: 512
	5420 Series	512
	5520 Series	1,024
	5720 Series	1,024

Table 16: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
IPv6 Neighbor table	5320 Series	8,000
	5420 Series	8,000 for 5420F Series models 16,000 for 5420M Series models
	5520 Series	16,000
	5720 Series	24,000 for 5720 MW Series models 32,000 for 5720 MXW Series models
IPv6 static entries in Source Binding Table	5320 Series	48-port models: 65 per VRF/ 256 per switch 16- and 24-port models: 256 per switch
	5420 Series	64 per VRF 256 per system
	5520 Series	128 per VRF 512 per system
	5720 Series	256
IPv6 static neighbor records	5320 Series	48-port models: 64 per VRF/256 per switch 16- and 24-port models: 256 per switch
	5420 Series	64 per VRF 256 per switch
	5520 Series	128 per VRF 512 per system
	5720 Series	128 per VRF 512 per system
IPv6 CLIP interfaces	5320 Series	64
	5420 Series	64
	5520 Series	64
	5720 Series	64
IPv6 static routes	5320 Series	48-port models: 501 16- and 24-port models: 500
	5420 Series	500
	5520 Series	1,000
	5720 Series	1,000

Table 16: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
IPv6 6in4 configured tunnels	5320 Series	32
	5420 Series	32
	5520 Series	64
	5720 Series	64
IPv6 DHCP Relay forwarding	5320 Series	248
	5420 Series	248
	5520 Series	256 per switch 10 per VRF
	5720 Series	512 per switch 10 per VRF
IPv6 BGP peers	5320 Series	8
	5420 Series	8
	5520 Series	16 Up to 8,000 IPv6 prefixes for BGPv6 peering
	5720 Series	256
IPv6 VRFs with iBGP	5320 Series	48-port models: 8 16- and 24-port models: 1
	5420 Series	8
	5520 Series	16
	5720 Series	16
BFD VRF instances	5320 Series	48-port models: 16 16- and 24- port models: 1
	5420 Series	16
	5520 Series	16
	5720 Series	16
BFD sessions per switch (IPv4/IPv6) with default values	5320 Series	48-port models: 16 16- and 24- port models: 1
	5420 Series	16
	5520 Series	16
	5720 Series	16
BFD sessions with Fabric Extend tunnels (IPv4)	5320 Series	48-port models: 16 16- and 24- port models: 1
	5420 Series	16
	5520 Series	16
	5720 Series	16

IP Interface Maximums for 5720 Series

The maximum number of IP interfaces for 5720 Series is based on the following formulas:

- If you disable the VRF scaling boot configuration flag:
 - # IP interfaces (1000 max) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + (# of RSMILT interfaces) + 2(if IP Shortcuts is enabled) + 3x(# of VRFs) = cannot exceed 1000
- If you enable the VRF scaling boot configuration flag:
 - # IP interfaces (max 1000) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + (# of RSMILT interfaces) + 2(if IP Shortcuts is enabled) + 3 = cannot exceed 1000

IP Interface Maximums for 5520 Series

The maximum number of IP interfaces for 5520 Series is based on the following formulas:

- If you disable the VRF scaling boot configuration flag:
 - # IP interfaces (500 max) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + (# of RSMILT interfaces) + 2(if IP Shortcuts is enabled) + 3x(# of VRFs) = cannot exceed 1000
- If you enable the VRF scaling boot configuration flag:
 - # IP interfaces (max 500) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + (# of RSMILT interfaces) + 2(if IP Shortcuts is enabled) + 3 = cannot exceed 1000

IP Interface Maximums for 5420 Series

The maximum number of IP interfaces for 5420 Series is based on the following formulas:

- If you disable the VRF scaling boot configuration flag:
 - # IP interfaces (248 max) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + (# of RSMILT interfaces) + 2(if IP Shortcuts is enabled) + 3x(# of VRFs) = cannot exceed 248
- If you enable the VRF scaling boot configuration flag:
 - # IP interfaces (max 248) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + (# of RSMILT interfaces) + 2(if IP Shortcuts is enabled) + 3 = cannot exceed 248

IP Interface Maximums for 5320 Series

The maximum number of IP interfaces for 5320 Series is based on the following formulas:

16- and 24-port models

- # IP interfaces (max 248) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + 2(if IP Shortcuts is enabled) + 3 = cannot exceed 248

48-port models

- If you disable the VRF scaling boot configuration flag:
 - # IP interfaces (248 max) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + 2(if IP Shortcuts is enabled) + 3x(# of VRFs) = cannot exceed 248

- If you enable the VRF scaling boot configuration flag:
 - # IP interfaces (max 248) + (# of VRRP IPv4 interfaces) + (# of VRRP IPv6 interfaces) + 2(if IP Shortcuts is enabled) + 3 = cannot exceed 248

Layer 3 Route Table Size

Table 17: Layer 3 Route Table Size Maximums

Attribute	Maximum number supported
IPv4 RIP routes	See Route Scaling on page 44.
IPv4 OSPF routes	
IPv4 BGP routes	
IPv4 SPB shortcut routes	
IPv4 SPB Layer 3 VSN routes	
IPv6 OSPFv3 routes - GRT only	
IPv6 SPB shortcut routes - GRT only	
IPv6 RIPng routes	

Route Scaling

The following table provides information on IPv4 and IPv6 route scaling. The route table is a shared hardware resource where IPv4 routes consume one entry and IPv6 routes with a prefix length less than 64 consume two entries.

The route scaling does not depend on the protocol itself, but rather the general system limitation in the following configuration modes:

- URPF check mode - Enable this boot configuration flag to support Unicast Reverse Path Forwarding check mode.
- IPv6 mode - Enable this boot configuration flag to support IPv6 routes with prefix-lengths greater than 64 bits. When the IPv6-mode is enabled, the maximum number of IPv4 routing table entries decreases. This flag does not apply to all hardware platforms.

**Note**

Only 5320-48P-8XE and 5320-48T-8XE support URPF mode.

Table 18: 5320 Series

URPF mode	IPv6 mode	5320 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
No	No	48-port models: 12K 16- and 24-port models: 8K	48-port models: 6K 16- and 24-port models: 4K	n/a
No	Yes	48-port models: 6K 16- and 24-port models: 4K	48-port models: 2K 16- and 24-port models: 2K	48-port models: 1.5K 16- and 24-port models: 1K
Yes	No	48-port models: 6K	48-port models: 2K	n/a
Yes	Yes	48-port models: 3K	48-port models: 1K	48-port models: 750

Note:
The total number of routes include local routes.
The stated numbers in the preceding rows are one-dimensional where the given number implies that only routes for that address family or type are present. For a given row in the table, the maximum scaling number is 'x' IPv4 routes OR 'y' ipv6 <= 64 routes (not a combination of both).

Table 19: 5420 Series

URPF mode	IPv6 mode	5420 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
No	No	12K	6K	n/a
No	Yes	6K	2K	1,500
Yes	No	6K	3K	n/a
Yes	Yes	3K	1K	750

Note:
The total number of routes include local routes.
The stated numbers in the preceding rows are one-dimensional where the given number implies that only routes for that address family or type are present. For a given row in the table, the maximum scaling number is 'x' IPv4 routes OR 'y' ipv6 <= 64 routes (not a combination of both).

Table 20: 5520 Series

URPF mode	IPv6 mode	5520 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
No	No	16K	8K	n/a
No	Yes	8K	4K	2K
Yes	No	8K	4K	n/a
Yes	Yes	4K	2K	1K

Note:

The total number of routes include local routes.

The stated numbers in the preceding rows are one-dimensional where the given number implies that only routes for that address family or type are present. For a given row in the table, the maximum scaling number is 'x' IPv4 routes OR 'y' ipv6 <= 64 routes (not a combination of both).

Table 21: 5720 Series

URPF mode	IPv6 mode	5720 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
No	No	5720 MW Series models: 16K 5720 MXW Series models: 24K	5720 MW Series models: 8K 5720 MXW Series models: 12K	n/a
No	Yes	5720 MW Series models: 8K 5720 MXW Series models: 12K	5720 MW Series models: 4K 5720 MXW Series models: 6K	5720 MW Series models: 2K 5720 MXW Series models: 3K
Yes	No	5720 MW Series models: 8K 5720 MXW Series models: 12K	5720 MW Series models: 4K 5720 MXW Series models: 6K	n/a
Yes	Yes	5720 MW Series models: 4K 5720 MXW Series models: 6K	5720 MW Series models: 2K 5720 MXW Series models: 3K	5720 MW Series models: 1K 5720 MXW Series models: 1.5K

Note:

The total number of routes include local routes.

The stated numbers in the preceding rows are one-dimensional where the given number implies that only routes for that address family or type are present. For a given row in the table, the maximum scaling number is 'x' IPv4 routes OR 'y' ipv6 <= 64 routes (not a combination of both).

IP Multicast

Table 22: IP Multicast Maximums

Attribute	Product	Maximum number supported
IGMP/MLD interfaces (IPv4/IPv6)	5320 Series	4,000/2,000
	5420 Series	4,000/2,000
	5520 Series	4,059
	5720 Series	4,059
PIM interfaces (IPv4/IPv6)	5320 Series	16 active
	5420 Series	16 active
	5520 Series	128 active
	5720 Series	128 active
PIM Neighbors (IPv4/IPv6) (GRT Only)	5320 Series	16
	5420 Series	16
	5520 Series	128
	5720 Series	128
PIM-SSM static channels (IPv4/IPv6)	5320 Series	512
	5420 Series	512
	5520 Series	4,000
	5720 Series	4,000
Multicast receivers/IGMP joins (IPv4/IPv6) (per switch)	5320 Series	6,000
	5420 Series	6,000
	5520 Series	6,000
	5720 Series	6,000
Total multicast routes (S,G,V) (IPv4/IPv6) (per switch)	5320 Series	48-port models: 4,000 16- and 24-port models: 2,000
	5420 Series	4,000
	5520 Series	4,000
	5720 Series	6,000
Total multicast routes (S,G,V) (IPv4) on an SPB-PIM Gateway configured switch	5320 Series	n/a
	5420 Series	n/a
	5520 Series	4,000
	5720 Series	n/a

Table 22: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
Static multicast routes (S,G,V) (IPv4/IPv6)	5320 Series	48-port models: 4,000 16- and 24-port models: 2,000
	5420 Series	4,000
	5520 Series	4,000
	5720 Series	6,000
Multicast enabled Layer 2 VSN (IPv4)	5320 Series	48-port models: 500 16- and 24-port models: 250
	5420 Series	500
	5520 Series	2,000
	5720 Series	2,000
Multicast enabled Layer 3 VSN (IPv4)	5320 Series	48-port models: 64 16- and 24-port models: 1
	5420 Series	64
	5520 Series	256 including mgmt VRF and GRT
	5720 Series	256
SPB-PIM Gateway controller S,Gs (source announcements) with MSDP (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	6,000
	5720 Series	n/a
SPB-PIM Gateway controllers per SPB fabric (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	5
	5720 Series	n/a
SPB-PIM Gateway nodes per SPB fabric (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	64
	5720 Series	n/a
SPB-PIM Gateway interfaces per BEB (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	64
	5720 Series	n/a

Table 22: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
PIM neighbors per SPB-PIM Gateway node (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	64
	5720 Series	n/a

Distributed Virtual Routing (DvR)

**Note**

Local hosts use ARP entries and remote hosts use host entries. For information on IP ARP scaling, see [IP Unicast](#) on page 36.

Table 23: DvR Maximums

Attribute	Product	Maximum number supported
Note: <ul style="list-style-type: none"> On the DvR leaf, you must enable the VRF scaling boot configuration flag if more than 24 VRFs are required in the DvR domain. Scaling of a VSP 4450 Series switch controls the scaling of the DvR domain it is in. For example, if a VSP 4450 Series switch is in a DvR domain with other platforms such as 5520 Series and 5420 Series, the scaling of the entire domain is limited to the scaling of the VSP 4450 Series. 		
DvR Virtual IP interfaces	5320 Series	48-port models: 248 16- and 24-port models: n/a
	5420 Series	247 with VIST 248 without VIST
	5520 Series	499 with vIST 500 without vIST
	5720 Series	999 with vIST 1000 without vIST
DvR domains per SPB fabric	5320 Series	16
	5420 Series	16
	5520 Series	16
	5720 Series	16

Table 23: DvR Maximums (continued)

Attribute	Product	Maximum number supported
Controller nodes per DvR domain with default route inject flag enabled Total number of Controllers per domain cannot exceed 8. Note: A DvR domain containing only Controller nodes and no Leaf nodes can have more than 8 Controllers per domain.	5320 Series	n/a
	5420 Series	n/a
	5520 Series	8
	5720 Series	8
Leaf nodes per DvR domain	5320 Series	250
	5420 Series	250
	5520 Series	250
	5720 Series	250
DvR enabled Layer 2 VSNs	5320 Series	48-port models: 248 16- and 24-port models: n/a
	5420 Series	247 with vIST 248 without vIST
	5520 Series	499 with vIST 500 without vIST
	5720 Series	999 with vIST 1000 without vIST
DvR host route scaling per DvR domain (scaling number includes local as well as foreign hosts of the Layer 2 VSN that are members of the domain) If DvR Layer 2 VSNs span DvR domains, and all DvR Controllers have an IP interface on the Layer 2 VSNs, then the DvR host scaling is network-wide, as DvR Controllers will consume as many host routes as there are hosts across all DvR domains.	5320 Series	48-port models: 16,000 16- and 24-port models: n/a
	5420 Series	16,000 for 5420F Series models 32,000 for 5420M Series models
	5520 Series	48,000
	5720 Series	64,000 for 5720MW Series models 96,000 for 5720MXW Series models

Filters, QoS, and Security

Table 24: Filters, QoS, and Security Maximums

Attribute	Product	Maximum number supported
For more information, see Filter Scaling on page 53.		
Total IPv4 Ingress rules/ACEs (Port/VLAN/InVSN based, Security/QoS filters)	5320 Series	48-port models: 3,072 16- and 24-port models: 1024
	5420 Series	2,048 Primary Bank 1,024 Secondary Bank
	5520 Series	1,024 Primary Bank 512 Secondary Bank
	5720 Series	
	5720MW Series models	Primary Bank: 3,072 Secondary Bank: 1,536
	5720MXW Series models	Primary Bank: 4,096 Secondary Bank: 2,048
Total IPv4 Egress rules/ACEs (Port based, Security filters)	5320 Series	48-port models: 400, 144 if you enable boot config flags ipv6-egress-filter or boot config flags macsec 16- and 24-port models: 190, 62 if you enable boot config flags ipv6-egress-filter or boot config flags macsec
	5420 Series	400 144 if you enable boot config flags ipv6-egress-filter or boot config flags macsec
	5520 Series	336 80 if you enable boot config flags ipv6- egress-filter
	5720 Series	5720MW Series models: 2,982 1,446 if you enable boot config flags ipv6- egress-filter 5720MXW Series models: 6,000 2,982 if you enable boot config flags ipv6- egress-filter

Table 24: Filters, QoS, and Security Maximums (continued)

Attribute	Product	Maximum number supported
Total IPv6 Ingress rules/ACEs (Port/VLAN/InVSN based, Security filters)	5320 Series	1,024
	5420 Series	512
	5520 Series	512
	5720 Series	1,536 for 5720MW Series models 2,048 for 5720MXW Series models:
Total IPv6 egress rules/ACEs (Port based, Security filters)	5320 Series	48-port models: 256, 0 with MACsec 16- and 24-port models: 128, 0 with MACsec
	5420 Series	256, 0 with MACsec
	5520 Series	256
	5720 Series	1,536 for 5720MW Series models 3,072 for 5720MXW Series models:
EAP (clients per port) Note: The total of EAP clients plus NEAP clients per port or per switch cannot exceed 8,192.	5320 Series	32
	5420 Series	32
	5520 Series	32
	5720 Series	32

Table 25: NEAP Maximums

Product	Max # supported	Details
5320 Series Note: The total of EAP clients plus NEAP clients per port or per switch cannot exceed 8,192. Note: Resources are shared with Switched UNI Endpoints.	800	boot config flags macsec: NO boot config flags spbm-node-scaling: NO Platform VLAN: N/A
	800	boot config flags macsec: YES boot config flags spbm-node-scaling: NO Platform VLAN: NO
	700	boot config flags macsec: YES boot config flags spbm-node-scaling: NO Platform VLAN: YES
	400	boot config flags macsec: N/A

Table 25: NEAP Maximums (continued)

Product	Max # supported	Details
		boot config flags spbm-node-scaling: YES Platform VLAN: N/A
5420 Series	800	boot config flags macsec: NO boot config flags spbm-node-scaling: NO Platform VLAN: N/A
	800	boot config flags macsec: YES boot config flags spbm-node-scaling: NO Platform VLAN: NO
	700	boot config flags macsec: YES boot config flags spbm-node-scaling: NO Platform VLAN: YES
	400	boot config flags macsec: N/A boot config flags spbm-node-scaling: YES Platform VLAN: N/A
5520 Series	4,900	N/A
5720 Series	8192	N/A

Filter Scaling

This section provides more details on filter scaling numbers for the universal hardware platforms.

5320 Series

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
 - 512 ACLs with 5 ACEs each that can hold either Security/QoS/both action types or
 - a combination based on the following rule: $(\text{num ACLs} + \text{num ACEs}) \leq 3072$

This maximum implies a VLAN member count of 1 for inVlan ACLs or a single I-SID for inVSN. The number of rules consumed by IPv4 inPort ACLs is not multiplied by the number of ports to which this ACL applies.

- Up to 1000 ACEs in a single ACL
- 512 IPv6 ingress ACLs (inPort):
 - 512 ACLs with 5 ACEs each that can hold either Security/QoS/both action types or
 - a combination based on the following rule: $(\text{num ACLs} + \text{num ACEs}) \leq 3072$

This maximum also implies a port member count of 1 for the inPort ACL. The number of rules consumed by IPv6 inPort ACL is multiplied by the number of ports to which this ACL applies.

- 1024 ingress ACEs: All ACEs can hold either Security/QoS/both action types

This maximum also implies a VLAN member count of 1 for an inVlan ACL.

- 400 egress ACEs

This maximum also implies a port member count of 1 for the outPort ACL.

5420 Series

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
 - 512 ACLs with 3 Primary Bank ACEs each OR
 - 512 ACLs with 1 Security Bank ACE each OR
 - a combination based on the following rule:
 - $((\text{num ACLs} + \text{num Primary Bank ACEs}) \leq 2048) \ \&\& \ ((\text{num ACLs} + \text{num Secondary Bank ACEs}) \leq 1024)$

This maximum implies a VLAN member count of 1 for inVlan ACLs or a single I-SID for inVSN. The number of rules consumed by IPv4 inPort ACLs is not multiplied by the number of ports to which this ACL applies.

- Up to 1000 ACEs in a single ACL
- 512 IPv6 ingress ACLs (inPort):
 - 512 ACLs with 1 ACE each OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num IPv6 ACEs} + \text{num IPv4 Secondary Bank ACEs}) \leq 1024$

This maximum also implies a port member count of 1 for the inPort ACL. The number of rules consumed by IPv6 inPort ACLs is multiplied by the number of ports to which this ACL applies.

- 3072 ingress ACEs:

Theoretical maximum of 1024 implies 1 ingress ACL with 512 Primary Bank ACEs and 512 Secondary Bank ACEs

- Ingress ACEs supported: $(2048 \text{ (Primary Bank)} - \# \text{ of ACLs}) + (1024 \text{ (Secondary Bank)} - \# \text{ of ACLs})$.

This maximum also implies a VLAN member count of 1 for an inVlan ACL.

- 400 egress ACEs:

Theoretical maximum of 400 implies 1 egress ACL with 400 ACEs

- Egress ACEs supported: $400 - \# \text{ of ACLs}$.

This maximum also implies a port member count of 1 for the outPort ACL.

5520 Series

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
 - 512 ACLs with 1 Primary ACE each OR
 - 256 ACLs with 1 Secondary ACE each OR
 - a combination based on the following rule:
 - $((\text{num ACLs} + \text{num Primary Bank ACEs}) \leq 1024) \ \&\& \ ((\text{num ACLs} + \text{num Secondary Bank ACEs}) \leq 512)$

This maximum implies a VLAN member count of 1 for inVlan ACLs or a single I-SID for inVSN. The number of rules consumed by IPv4 inPort ACLs is not multiplied by the number of ports to which this ACL applies.

- Up to 1000 ACEs in a single ACL
- 512 IPv6 ingress ACLs (inPort):
 - 512 ACLs with 1 ACE each OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num ACEs} + \text{num IPv4 Security Bank ACEs}) \leq 512$

The number of rules consumed by IPv6 ingress ACLs inPort ACLs is multiplied by the number of ports to which this ACL applies.

- 124 egress ACLs (outPort only):
 - 124 ACLs with 1 ACE each (one of these ACLs can have 2 ACEs) OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num ACEs}) \leq 248$

This maximum implies a port member count of 1 for outPort ACLs.

- 1536 ingress ACEs:
 - Ingress ACEs supported: $(1024 \text{ (Primary Bank)} - \# \text{ of ACLs}) + (512 \text{ (Secondary Bank)} - \# \text{ of ACLs})$.
- 247 egress ACEs:
 - Egress ACEs supported: $248 - \# \text{ of ACLs}$.

This maximum also implies a port member count of 1 for the outPort ACL.

5720-24MW and 5720-48MW

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
 - 512 ACLs with 5 Primary Bank ACEs each OR
 - 512 ACLs with 2 Secondary Bank ACEs each OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num Primary Bank ACEs}) \leq 3072) \ \&\& \ ((\text{num ACLs} + \text{num Security Bank ACEs}) \leq 1536)$

This maximum implies a VLAN member count of 1 for inVlan ACLs or a single I-SID for inVSN.

The number of rules consumed by IPv4 inPort ACLs is not multiplied by the number of ports to which this ACL applies.

- You can configure up to 1000 ACEs in a single ACL.
- 512 IPv6 ingress ACLs (inPort):
 - 512 ACLs with 2 ACEs each OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num ACEs} + \text{num of IPv4 Security Bank ACEs}) \leq 1536$

The number of rules consumed by IPv6 inPort ACLs is multiplied by the number of ports to which this ACL applies.

- 256 egress ACLs (outPort only):
 - 1 OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num ACEs}) \leq 2982$
- 4608 ingress ACEs

Ingress ACEs supported: $(3072 \text{ Primary Bank} - \text{num ACLs}) + (1536 \text{ Secondary Bank} - \text{num ACEs})$

- 2982 egress ACEs

Egress ACEs supported: $2982 - \text{num ACLs}$

5720-24MXW and 5720-48MXW

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
 - 512 ACLs with 7 Primary Bank ACEs each OR
 - 512 ACLs with 3 Secondary Bank ACEs each OR
 - a combination based on the following rule:
 - $((\text{num ACLs} + \text{num Primary Bank ACEs}) \leq 4096) \ \&\& \ ((\text{num ACLs} + \text{num Security Bank ACEs}) \leq 2048)$

This maximum implies a VLAN member count of 1 for inVlan ACLs or a single I-SID for inVSN.

The number of rules consumed by IPv4 inPort ACLs is not multiplied by the number of ports to which this ACL applies.

- You can configure up to 1000 ACEs in a single ACL.
- 512 IPv6 ingress ACLs (inPort):
 - 512 ACLs with 3 ACEs each OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num ACEs} + \text{num of IPv4 Security Bank ACEs}) \leq 2048$

The number of rules consumed by IPv6 inPort ACLs is multiplied by the number of ports to which this ACL applies.

- 256 egress ACLs (outPort only):
 - 1 OR

- a combination based on the following rule:
 - (num ACLs + num ACEs) <=6000
- 6144 ingress ACEs

Ingress ACEs supported: (4096 Primary Bank - num ACLs) + (2048 Secondary Bank - num ACEs)

- 6000 egress ACEs

Egress ACEs supported: 6000 - num ACLs

Routed Private VLANs/E-TREES Scaling

The number of private VLANs that you configure with an IP address influences the IPv4 Egress ACE count.

The following table lists scaling limits for Routed Private VLANs/E-TREES. Limits are not enforced; either number of private VLANs or number of private VLAN trunk ports can go beyond the recommended values.

Table 26: Routed Private VLANs/E-TREES Maximums

	Private VLAN trunk ports	Routed PVLANS/E-TREES	IPv4 Egress ACE rules available (No IPv6 egress filter bootflag enabled)	IPv4 Egress ACE rules available (With IPv6 egress filter bootflag enabled)
5320-48T-8XE 5320-48P-8XE	4	10	349	93
5320-16P-4XE 5320-16P-4XE-DC 5320-24P-8XE 5320-24T-8XE	4	10	139	11
5420 Series	4	10	349	93
5520 Series	4	10	285	29
5720-24MW 5720-48MW	4	100	2499	999
5720-24MXW 5720-48MXW	4	100	5499	2499

Use the **show io resources filter** command to verify remaining resources. This command displays the following information:

- resources consumed by Routed Private VLANs
- free entries available for either IPv4 Egress ACEs or private VLANs

The following example output displays resource usage on a 5320 Series switch with one Routed Private VLAN and one outPort ACL.

```
Switch:1>show io resources filter
```

```
=====
FILTER TABLE
=====
```

```

-----
ACL Filter Resource Manager stats
-----
BCM CAP Group: | ICAP_SEC_QOS | ICAP_IPv6 | ECAP_SEC | ECAP_IPv6
Group Mode: | Double | Double | Double | Double
-----
Total Entries: | 1024 | 1024 | 247 | 128
Free Entries: | 1024 | 1024 | 243 | 128
In Use: | 0 | 0 | 4 | 0
Filter table:
-----
ACL | |Port/Vlan| Sec | QoS | All |
ID | Flags | Members | ACE's | ACE's | ACE's | Type
-----
1 | 00002008 | 1 | 0 | 0 | 1 | outPort, non-IPv6
-----

Filter resources used by other features:
-----
Feature | Type | Number of entries |
-----
PVlan | ECAP | 2 |
-----

```

OAM and Diagnostics

Table 27: OAM and Diagnostics Maximums

Attribute	Product	Maximum number supported
EDM sessions	5320 Series	5
	5420 Series	5
	5520 Series	5
	5720 Series	5
FTP sessions (IPv4/IPv6)	5320 Series	8 total (4 for IPv4 and 4 for IPv6)
	5420 Series	8 total (4 for IPv4 and 4 for IPv6)
	5520 Series	8 total (4 for IPv4 and 4 for IPv6)
	5720 Series	8 total (4 for IPv4 and 4 for IPv6)
SSH sessions (IPv4/IPv6)	5320 Series	8 total (any combination of IPv4 and IPv6)
	5420 Series	8 total (any combination of IPv4 and IPv6)
	5520 Series	8 total (any combination of IPv4 and IPv6)
	5720 Series	8 total (any combination of IPv4 and IPv6)

Table 27: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
Telnet sessions (IPv4/IPv6)	5320 Series	16 total (8 for IPv4 and 8 for IPv6)
	5420 Series	16 total (8 for IPv4 and 8 for IPv6)
	5520 Series	16 total (8 for IPv4 and 8 for IPv6)
	5720 Series	16 total (8 for IPv4 and 8 for IPv6)
TFTP sessions (IPv4/IPv6)	5320 Series	2 total (any combination of IPv4 and IPv6)
	5420 Series	2 total (any combination of IPv4 and IPv6)
	5520 Series	2 total (any combination of IPv4 and IPv6)
	5720 Series	2 total (any combination of IPv4 and IPv6)
Mirrored ports (source)	5320 Series	48-port models: 56 24-port models: 32 16-port models: 20
	5420 Series	56
	5520 Series	48-port models: 47 (up to 58 with channelization) 24-port models: 23 (up to 34 with channelization)
	5720 Series	64
Mirroring ports (destination)	5320 Series	4
	5420 Series	4
	5520 Series	4
	5720 Series	4

Table 27: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
Fabric RSPAN Port mirror instances per switch (Ingress only)	5320 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	5420 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	5520 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	5720 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
Fabric RSPAN Flow mirror instances per switch (Ingress only)	5320 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	5420 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	5520 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	5720 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
Fabric RSPAN Monitoring I-SIDs (network value)	5320 Series	48-port models: 500 Monitoring I-SIDs across SPB network 16 and 24-port models: 250 Monitoring I-SIDs across SPB network
	5420 Series	500 Monitoring I-SIDs across SPB network
	5520 Series	1,000 Monitoring I-SIDs across SPB network
	5720 Series	1,000 Monitoring I-SIDs across SPB network

Table 27: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
sFlow sampling limit	5320 Series	3,100 samples per second
	5420 Series	3,100 samples per second
	5520 Series	3,100 samples per second
	5720 Series	3,100 samples per second
IPFIX flows	5320 Series	48-port models: 9,000 16- and 24-port models: n/a
	5420 Series	9,000
	5520 Series	36,863
	5720 Series	5720MW models: 32,000 5720MXW models: 256,000
Application Telemetry host monitoring - maximum number of monitored hosts Note: These resources are shared with the IPv4 Filter Ingress rules/ACEs.	5320 Series	382 hosts
	5420 Series	382 hosts
	5520 Series	382 hosts
	5720 Series	382 hosts

Fabric Scaling

This section lists the fabric scaling information.

Table 28: Fabric Maximums

Attribute	Product	Maximum number supported (with and without vIST)
Number of SPB IS-IS areas	5320 Series	1
	5420 Series	1
	5520 Series	1
	5720 Series	1
Number of B-VIDs	5320 Series	2
	5420 Series	2
	5520 Series	2
	5720 Series	2

Table 28: Fabric Maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
Maximum number of Physical and Logical (Fabric Extend) NNI interfaces/adjacencies	5320 Series	64
	5420 Series	50
	5520 Series	128
	5720 Series	128
SPBM enabled nodes per area (BEB + BCB)	5320 Series	350 without boot config flags spbm-node-scaling 500 with boot config flags spbm-node-scaling
	5420 Series	350 without boot config flags spbm-node-scaling 500 with boot config flags spbm-node-scaling
	5520 Series	800
	5720 Series	1000
Number of BEBs not part of vIST clusters this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI)	5320 Series	350 without boot config flags spbm-node-scaling 500 with boot config flags spbm-node-scaling
	5420 Series	350 without boot config flags spbm-node-scaling 500 with boot config flags spbm-node-scaling
	5520 Series	800
	5720 Series	2,000
Number of BEBs that are part of a vIST cluster this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI)	5320 Series	300
	5420 Series	300
	5520 Series	800
	5720 Series	1,000

Table 28: Fabric Maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
I-SIDs supported (local UNI present on device)	5320 Series	See Number of I-SIDs supported
	5420 Series	See Number of I-SIDs supported
	5520 Series	See Number of I-SIDs supported
	5720 Series	See Number of I-SIDs supported
Maximum number of Layer 2 VSNs per switch (local UNI present on device)	5320 Series	48-port models: 500 16- and 24-port models: 250
	5420 Series	500
	5520 Series	3,580
	5720 Series	4,000
Maximum number of Transparent Port UNIs per switch	5320 Series	48-port models: 53 24-port models: 29 16- models: 20
	5420 Series	56
	5520 Series	48-port models: 48 24-port models: 24
	5720 Series	60
Maximum number of Layer 2 E-Tree/PVLAN UNIs per switch	5320 Series	48-port models: 50 16-port and 24-port models: 20
	5420 Series	100
	5520 Series	200
	5720 Series	100
Maximum number of routed PVLANS/E-Trees	5320 Series	10
	5420 Series	10
	5520 Series	10
	5720 Series	100

Table 28: Fabric Maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
Maximum number of Switched UNI Endpoints (C-VID or untagged port bindings)	5320 Series	See Maximum Number of Switched UNI Endpoints (C-VID or Untagged Port Bindings) on page 65.
	5420 Series	See Maximum Number of Switched UNI Endpoints (C-VID or Untagged Port Bindings) on page 65.
	5520 Series	See Maximum Number of Switched UNI Endpoints (C-VID or Untagged Port Bindings) on page 65.
	5720 Series	See Maximum Number of Switched UNI Endpoints (C-VID or Untagged Port Bindings) on page 65.
Maximum number of Layer 3 VSNs per switch See VRF Scaling on page 72.	5320 Series	48-port models: 64 16- and 24-port models: 1 local VRF and 23 remote accepted I-SIDs
	5420 Series	64
	5520 Series	256 including mgmt VRF and GRT
	5720 Series	256
Maximum number of SPB Layer 2 multicast Data I-SIDs	5320 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67
	5420 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67
	5520 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67
	5720 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67

Table 28: Fabric Maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
Maximum number of SPB Layer 3 multicast Data I-SIDs	5320 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67
	5420 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67
	5520 Series	Maximum 4,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached. See Maximum Number of SPB Multicast Data I-SIDs on page 67
	5720 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 67
Maximum number of FA ISID/VLAN assignments per port	5320 Series	94
	5420 Series	94
	5520 Series	94
	5720 Series	94
Maximum number of IP multicast S,Gs when operating as a BCB (intra-area)	5320 Series	16,000
	5420 Series	16,000
	5520 Series	16,000
	5720 Series	50,000

Maximum Number of Switched UNI Endpoints (C-VID or Untagged Port Bindings)

The number of Switched UNI endpoints supported varies based on other feature configuration.

The following table provides the maximum number supported with and without vIST.

Product	Total Number of Switched UNI Endpoints	Feature configuration
5320 Series Note: Resources are shared with NEAP clients.	800	boot config flags macsec: NO boot config flags spbm-node-scaling: NO Platform VLAN present: N/A
	800	boot config flags macsec: YES boot config flags spbm-node-scaling: NO Platform VLAN present: NO
	700	boot config flags macsec: YES boot config flags spbm-node-scaling: NO Platform VLAN present: YES
	400	boot config flags macsec: N/A boot config flags spbm-node-scaling: YES Platform VLAN present: N/A
5420 Series	800	boot config flags macsec: NO
	700	VLAN-based with boot config flags macsec
	400	boot config flags spbm-node-scaling: YES
5520 Series	4,900	N/A
5720 Series	12,000	N/A

Maximum Number of SPB Multicast Data I-SIDs

The number of I-SIDs supported varies for Layer 2 and Layer 3 ingress and egress BEBs.

Attribute		Product	Maximum number supported (with and without vIST)
Maximum number of Layer 2 multicast Data I-SIDs Note: Overall limits across all Layer 2 VSNs	On Ingress BEB: Dynamic and Static originated Data I-SIDs	5320 Series	16- and 24-port models: 250 48- port models: 500
		5420 Series	500
		5520 Series	3,580
		5720 Series	4,000
	On Egress BEB: Static Data I-SIDs Terminated	5320 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5420 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5520 Series	4,000
		5720 Series	6,000
	On Egress BEB: Dynamic data I-SIDs + originating BEB pairs terminated	5320 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5420 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags

Attribute		Product	Maximum number supported (with and without vIST)
			spbm-node-scaling
		5520 Series	4,000
		5720 Series	6,000

Attribute		Product	Maximum number supported (with and without vIST)
Maximum number of Layer 3 multicast Data I-SIDs Note: Overall limits across all Layer 3 VSNs/GRT	On Ingress BEB: Dynamic and Static originated Data I-SIDs	5320 Series	16- and 24-port models: 250 48- port models: 500
		5420 Series	500
		5520 Series	3,580
		5720 Series	4,000
	On Egress BEB: Static Data I-SIDs Terminated	5320 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5420 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5520 Series	4,000
		5720 Series	6,000
	On Egress BEB: Dynamic data I-SIDs terminated, each I-SID counted after each originating BEB	5320 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5420 Series	800 without boot config flags spbm-node-scaling 1,200 with boot config flags spbm-node-scaling
		5520 Series	4,000
		5720 Series	6,000

Number of I-SIDs Supported for the Number of Configured IS-IS Interfaces and Adjacencies (NNIs)

The number of I-SIDs supported depends on the number of IS-IS interfaces and adjacencies (NNIs) configured.

The following table shows the number of UNI I-SIDs supported per BEB. UNI I-SIDs are used for Layer 2 VSN, Layer 3 VSN, Transparent-UNI, E-Tree, Switched-UNI and S, G for Multicast.

Number of IS-IS interfaces (NNIs)	Product	I-SIDs with vIST configured on the platform	I-SIDs without vIST configured on the platform
4	5320 Series	n/a	500
	5420 Series	564	564
	5520 Series	4,000	4,000
	5720 Series	4,000	4,000
6	5320 Series	n/a	500
	5420 Series	564	564
	5520 Series	3,500	4,000
	5720 Series	3,500	4,000
10	5320 Series	n/a	500
	5420 Series	564	564
	5520 Series	2,900	4,000
	5720 Series	2,900	4,000
20	5320 Series	n/a	500
	5420 Series	564	564
	5520 Series	2,000	4,000
	5720 Series	2,000	4,000
48	5320 Series	n/a	500
	5420 Series	564	564
	5520 Series	1,000	2,000
	5720 Series	1,000	2,000
72	5320 Series	n/a	n/a
	5420 Series	n/a	n/a
	5520 Series	750	1,500
	5720 Series	750	1,500
100	5320 Series	n/a	n/a
	5420 Series	n/a	n/a
	5520 Series	550	1,100
	5720 Series	550	1,100

Number of IS-IS interfaces (NNIs)	Product	I-SIDs with vIST configured on the platform	I-SIDs without vIST configured on the platform
128	5320 Series	n/a	n/a
	5420 Series	n/a	n/a
	5520 Series	450	900
	5720 Series	450	900
250	5320 Series	n/a	n/a
	5420 Series	n/a	n/a
	5520 Series	n/a	n/a
	5720 Series	n/a	n/a

Interoperability Considerations for IS-IS External Metric

BEBs running VOSS 5.0 can advertise routes into IS-IS with the metric type as external. They can also correctly interpret route advertisements with metric type external received via IS-IS. In an SPB network with a mix of products running different versions of software releases, you must take care to ensure that turning on the ability to use metric-type external does not cause unintended loss of connectivity.

Note the following before turning on IS-IS external metric if the SPB network has switches running a release prior to VOSS 5.0:

- There are no special release or product type implications if the switch does not have IP Shortcuts or Layer 3 VSN enabled. For example, this applies to Layer 2 only BEBs and BCBs.
- There are no special release or product type implications if the Layer 3 VSN in which routes are being advertised with a metric-type of external is not configured on the switch.
- If a switch running a VOSS release that is prior to VOSS 5.0 but VOSS 4.2.1 or later, it will treat all IS-IS routes as having metric-type internal, regardless of the metric-type (internal or external) used by the advertising BEB in its route advertisement.
- Switches running VSP 9000 Series release 4.1.0.0 or later will treat all IS-IS routes as having metric-type internal, regardless of the metric-type (internal or external) used by the advertising BEB in its route advertisement.
- Switches running VOSS releases prior to 4.2.1.0 might not correctly install IS-IS routes in a Layer 3 VSN if any routes advertised with metric-type external are advertised in that Layer 3 VSN by other BEBs in the network. Layer 3 VSNs in which there are no routes with an external metric-type will not be impacted. Similar note applies to the GRT.
- Switches running VSP 9000 Series releases prior to 4.1.0.0 might not correctly install IS-IS routes in a Layer 3 VSN if any routes advertised with metric-type external are advertised in that Layer 3 VSN by other BEBs in the network. Layer 3 VSNs in which there are no routes with an external metric-type will not be impacted. Similar note applies to GRT.
- Switches running any ERS 8800 release might not correctly install IS-IS routes in a Layer 3 VSN if any routes advertised with metric-type external are advertised in that Layer 3 VSN by other BEBs in the network. Layer 3 VSNs in which there are no routes with an external metric-type will not be impacted. Similar note applies to GRT.

Recommendations

This section provides recommendations that affect feature configuration.

Pay special attention to the expected scaling of routes in the network and the number of OSPF neighbors in a single VRF when you select configuration values for the **isis 11-hellointerval** and **isis 11-hello-multiplier** commands on IS-IS interfaces. The default values for these commands work well for most networks, including those using moderately-scaled routes.

The default values work well for 16,000 routes and 64 OSPF neighbors in a single VRF. However, in highly-scaled networks, you might need to configure higher values for these commands.

For example, if the total number of non IS-IS routes on a given BEB exceeds 16,000 in combination with approximately 128 OSPF neighbors in a single VRF, you should configure a value of 12 for **isis 11-hellomultiplier**, instead of using the default value of 3.

VRF Scaling

By default, the system reserves VLAN IDs 4060 to 4094 for internal use.

If you enable both the VRF scaling and the SPBM mode boot configuration flags, the system reserves additional VLAN IDs (3500 to 3998) for internal use.

By default, VRF scaling is disabled and SPBM mode is enabled. When VRF scaling is disabled, you can have a maximum of 24 VRFs.

On the 5320 Series, only 5320-48P-8XE and 5320-48T-8XE support more than one VRF with IP configuration. The **boot config flag vrf-scaling** command does not apply to other 5320 Series models.



Important Notices

[ExtremeCloud IQ Support on page 73](#)

[Compatibility with ExtremeCloud IQ - Site Engine on page 73](#)

[Feature-Based Licensing on page 74](#)

[Memory Usage on page 74](#)

Unless specifically stated otherwise, the notices in this section apply to all platforms.

ExtremeCloud IQ Support

ExtremeCloud™ IQ provides cloud-managed networking, and delivers unified, full-stack management of wireless access points, switches, and routers. It enables onboarding, configuration, monitoring, troubleshooting, reporting, and more. Using innovative machine learning and artificial intelligence technologies, ExtremeCloud IQ analyzes and interprets millions of network and user data points, from the network edge to the data center, to power actionable business and IT insights, and to deliver new levels of network automation and intelligence.

ExtremeCloud IQ supports the following platforms:

- 5320 Series
- 5420 Series
- 5520 Series
- 5720 Series

For the most current information on switches supported by ExtremeCloud IQ, see [ExtremeCloud™ IQ Learning What's New](#).

The switch supports a zero touch connection to ExtremeCloud IQ. Zero touch deployment is used to deploy and configure a switch using ExtremeCloud IQ.

The switch software integrates with ExtremeCloud IQ using IQAgent.

For more information, see [Fabric Engine User Guide](#).

For more information about ExtremeCloud IQ, go to <https://www.extremenetworks.com/support/documentation/extremecloud-iq/>.

Compatibility with ExtremeCloud IQ - Site Engine

To understand which versions of ExtremeCloud IQ - Site Engine are compatible with this Network Operating System release on different hardware platforms, see [Extended Firmware Support](#).

Feature-Based Licensing

The following table provides information on the feature-licensing models available. For more information about licensing including feature inclusion, order codes, and how to load a license file, see [Fabric Engine User Guide](#).

Table 29: License models

Product	License model
5320 Series 5420 Series 5520 Series 5720 Series	Support a perpetual licensing model that includes Base, Premier, and MACsec licenses. Premier and MACsec licenses enable advanced features not available in the Base License. Because the hardware supports more than one Network Operating System (NOS) personality, it uses a licensing scheme that is NOS agnostic. Note: 5320 Series supports a 4-port and an 8-port 10G Port license.

Memory Usage

These switches intentionally reboot when memory usage on the switch reaches 95%.



Known Issues and Restrictions

[Known Issues on page 75](#)

[Restrictions and Expected Behaviors on page 97](#)

This section details the known issues and restrictions found in this release. Where appropriate, use the workarounds provided.

Known Issues

This section identifies the known issues in this release.

Known Issues for 8.9

Issue number	Description	Workaround
	HTTPS connection fails for CA-signed certificate with certificate inadequate type error on FF.	Ensure End-Entity, Intermediate CA and Root CA certificates are all SHA256 based and RSA2048 key signed, and Extended key usage field is set to TLS webservice Auth only for subject and root. For intermediate, it must be set with other required bits to avoid this issue. Add the root, intermediate CAs in the trust store of the browser for accessing the EDM with HTTPS.
VOSS-1265	On the port that is removed from a T-UNI LACP MLT, non T-UNI configuration is blocked as a result of T-UNI consistency checks.	When a port is removed from a T-UNI LACP MLT, the LACP key of the port must be set to default.
VOSS-1278	SLA Mon tests fail (between 2% and 8% failure) between devices when you have too many agents involved with scaled configurations.	This happens only in a scaled scenario with more than seven agents, otherwise the failure does not occur. The acceptable failure percentage is 5%, but you could see failures of up to 8%.

Issue number	Description	Workaround
VOSS-1280	The following error message occurs when performing shutdown/no-shutdown commands continuously: IO1 [05/02/14 06:59:55.178:UTC] 0x0011c52500000000 GlobalRouter COP-SW ERROR vsp4kTxEnable Error changing TX disable for SFP module: 24, code: -8	None. When this issue occurs, the port in question can go down, then performs a shutdown/no-shutdown of the port to bring it up and resumes operation.
VOSS-1285	CAKs are not cleared after setting the device to factory-default.	None. Currently this is the default behavior and does not affect functionality of the MACsec feature.
VOSS-1288	Shutting down the T1 link from one end of the link does not shut down the link at the remote end. You could experience traffic loss if the remote side of the link is not shut down.	This issue occurs only when a T1 SFP link from one end is shutdown. Enable a dynamic link layer protocol such as LACP or VLACP on both ends to shut the remote end down too. As an alternative, administratively disable both ends of the T1 SFP link to avoid the impact.
VOSS-1289	On a MACsec-enabled port, you can see delayed packets when the MACsec port is kept running for more than 12 hours. This delayed packet counter can also increment when there is complete reordering of packets so that the application might receive a slow response. But in this second case, it is a marginal increase in the packet count, which occurs due to PN mismatch sometimes only during Key expiry, and does not induce any latency.	None.
VOSS-1309	You cannot use EDM to issue ping or traceroute commands for IPv6 addresses.	Use CLI to initiate ping and traceroute commands.
VOSS-1310	You cannot use EDM to issue ping or traceroute commands for IPv4 addresses.	Use CLI to initiate ping and traceroute commands.
VOSS-1312	On the VSP 8400 Series 40-gigabit ports, the small metallic fingers that surround the ports are fragile and can bend out of shape during removal and insertion of the transceivers. When the fingers are bent, they prevent the insertion of the QSFP+ transceiver.	Insert the QSFP+ carefully. If the port becomes damaged, it needs to be repaired.

Issue number	Description	Workaround
VOSS-1335	<p>In an IGMP snoop environment, after dynamically downgrading the IGMP version to version 2 (v2), when you revert back to version 3 (v3), the following is observed:</p> <ul style="list-style-type: none"> • The multicast traffic does not flow. • The sender entries are not learned on the local sender switch. • The Indiscard packet count is incremented on the show int gig error statistics command. 	Use a v3 interface as querier in a LAN segment that has snoop-enabled v2 and v3 interfaces.
VOSS-1344	In EDM, you cannot select multiple 40 gigabit ports or a range of ports that includes 40 gigabit ports to graph or edit. You need to select them and edit them individually.	None.
VOSS-1349	On EDM, the port LED for channelized ports only shows the status of sub-port #1, but not the rest of the sub-ports. When you remove sub-port #1, and at least one other sub-port is active and online, the LED color changes to amber, when it should be green because at least one other sub-ports is active and online. The LED only shows the status of sub-port #1.	None.
VOSS-1354	An intermittent link-flap issue can occur in the following circumstance for the copper ports. If you use a crossover cable and disable auto-negotiation, the port operates at 100 Mbps. A link flap issue can occur intermittently and link flap detect will shut down the port.	Administratively shutdown, and then re-enable the port. Use auto-negotiation. Disabling auto-negotiation on these ports is not a recommended configuration.
VOSS-1358	Traffic is forwarded to IGMP v2 SSM group, even after you delete the IGMP SSM-map entry for the group.	If you perform the delete action first, you can recreate the SSM-map record, and then disable the SSM-map record. The disabled SSM-map record causes the receiver to timeout because any subsequent membership reports that arrive and match the disabled SSM-map record are dropped. You can delete the SSM-map record after the receivers time out.
VOSS-1359	The 4 byte AS confederation identifier and peers configuration are not retained across a reboot. This problem occurs when 4 Byte AS is enabled with confederation.	Reconfigure the 4 byte AS confederation identifier and peers on the device, and reboot.

Issue number	Description	Workaround
VOSS-1360	<p>After you enable enhanced secure mode, and log in for the first time, the system prompts you to enter a new password. If you do not meet the minimum password requirements, the system displays the following message: Password should contain a minimum of 2 upper and lowercase letters, 2 numbers and 2 special characters like !@#\$%^&*(). Password change aborted. Enter the New password:</p> <p>The system output message does not display the actual minimum password requirements you need to meet, which are configured on your system. The output message is an example of what the requirements need to meet. The actual minimum password requirements you need to meet are configured on your system by the administrator.</p>	None.
VOSS-1367	The configuration file always includes the router ospf entry regardless of whether OSPF is configured. This line does not perform any configuration and has no impact on the running software.	None.
VOSS-1368	When you use Telnet or SSH to connect to the switch, it can take up to 60 seconds for the log in prompt to appear. However, this situation is very unlikely to happen, and it does not appear in a standard normal operational network.	Do not provision DNS servers on a switch to avoid this issue altogether.
VOSS-1370	If you configure egress mirroring on NNI ports, you do not see the MAC-in-MAC header on captured packets.	Use an Rx mirror on the other end of the link to see the packets.
VOSS-1371	A large number of IPv6 VRRP VR instances on the same VLAN can cause high CPU utilization.	Do not create more than 10 IPv6 VRRP VRs on a single VLAN.
VOSS-1389	If you disable IPv6 on one RSMLT peer, the switch can intermittently display COP-SW ERROR and RCIP6 ERROR error messages. This issue has no impact.	None.
VOSS-1390	If you delete the SPBM configuration and re-configure SPBM using the same nickname but a different IS-IS system ID without rebooting, the switch displays an error message.	Reboot the switch after you delete the SPBM configuration.
VOSS-1403	EDM displays the user name as Admin, even though you log in using a different user name.	None.

Issue number	Description	Workaround
VOSS-1406	When you re-enable insecure protocols in the CLI SSH secure mode, the switch does not display a warning message.	None.
VOSS-1418	EDM displays the IGMP group entry that is learned on a vIST MLT port as TX-NNI.	Use CLI to view the IGMP group entry learned on a vIST MLT port.
VOSS-1428	When port-lock is enabled on the port and re-authentication on the EAP client fails, the port is removed from the RADIUS-assigned VLAN. This adds the port to the default VLAN and displays an error message. This issue has no impact.	The error message is incorrect and can be ignored.
VOSS-1433	When you manually enable or disable IS-IS on 40 Gbps ports with CR4 direct attach cables (DAC), the port bounces one time.	Configure IS-IS during the maintenance period. Bring the port down, configure the port and then bring the port up.
VOSS-1438	In a rare scenario in Simplified vIST configuration when vIST state is toggled immediately followed by vIST MLT ports are toggled, one of the MLT ports will go into blocking state resulting in failure to process data packets hashing to that link.	Before enabling vIST state ensure all vIST MLT ports are shut and re-enabled after vIST is enabled on the DUT.
VOSS-1440 VOSS-1441	When you configure a scaled Layer 3 VSN (24 Layer 3 VSN instances), route leaking from GRT to VRF on the local DUT does not happen. The switch displays an incorrect error message: <code>Only 24 Layer 3 VSNs can be configured.</code>	None.
VOSS-1463 VOSS-1471	When you use Fabric Extend over IP (FE-IP) and Fabric Extend over Layer 2 VLAN (FE-VID) solution, if you change the ingress and egress .1p map, packets cannot follow correct internal QoS queues for FE tunnel to FE tunnel, or FE tunnel to regular NNI traffic.	Do not change the default ingress and egress .1p maps when using Fabric Extend. With default ingress and egress .1p maps, packets follow the correct internal QoS when using the Fabric Extend feature.
VOSS-1473	If the I-SID associated with a Switched UNI or Fabric Attach port does not have a platform VLAN association and you disable Layer 2 Trusted, then the non IP traffic coming from that port does not take the port QoS and still uses the .1p priority in the packet.	None.
VOSS-1530	If you improperly close an SSH session, the session structure information does not clear and the client can stop functioning.	Disable and enable SSH.
VOSS-1584	The <code>show debug-file all</code> command is missing.	None.

Issue number	Description	Workaround
VOSS-1585	The system does not generate a log message, either in the log file or on screen, when you run the flight-recorder command.	None.
VOSS-1608	If you use an ERS 4850 FA Proxy with a VOSS or Fabric Engine FA Server, a mismatch can exist in the show output for tagged management traffic. The ERS device always sends traffic as tagged. The VOSS or Fabric Engine FA Server can send both tagged and untagged. For untagged, the VOSS and Fabric Engine FA Servers send VLAN ID 4095 in the management VLAN field of the FA element TLV. The ERS device does not recognize this VLAN ID and so still reports the traffic as tagged.	There is no functional impact.
VOSS-1706	EAPOL: Untagged traffic is not honoring the port QOS for Layer 2 trusted/ Layer 3 untrusted. This issue is only seen on EAPOL-enabled ports.	None.
VOSS-2014	IPv6 MLD Group is learned for Link-Local Scope Multicast Addresses. This displays additional entries in the Multicast routing tables.	None.

Issue number	Description	Workaround
VOSS-2033	<p>The following error messages appear when you use the shutdown and no shutdown commands on the MLT interface with ECMP and BGP+ enabled:</p> <pre> CP1 [01/23/16 11:10:16.474:UTC] 0x00108628 00000000 GlobalRouter RCIP6 ERROR rcIpReplaceRouteNotifyIpv6:FA IL ReplaceTunnelRec conn_id 2 CP1 [12/09/15 12:27:02.203:UTC] 0x00108649 00000000 GlobalRouter RCIP6 ERROR ifyRpcOutDelFibEntry: del FIB of Ipv6Route failed with 0: ipv6addr: 201:6:604:0:0:0:0:0, mask: 96, nh: 0:0:0:0:0:0:0:0 cid 6657 owner BGP CP1 [12/09/15 12:20:30.302:UTC] 0x00108649 00000000 GlobalRouter RCIP6 ERROR ifyRpcOutDelFibEntry: del FIB of Ipv6Route failed with 0: ipv6addr: 210:6:782:0:0:0:0:0, mask: 96, nh: fe80:0:0:0:b2ad:aaff:fe55:508 8 cid 2361 owner OSPF </pre>	Disable the alternate path.
VOSS-2036	IPsec statistics for the management interface do not increment for inESPFailures or InAHFailures.	None.
VOSS-2117	If you configure static IGMP receivers on an IGMPv3 interface and a dynamic join and leave are received on that device from the same destination VLAN or egress point, the device stops forwarding traffic to the static receiver group after the dynamic leave is processed on the device. The end result is that the IGMP static groups still exist on the device but traffic is not forwarded.	Disable and re-enable IGMP Snooping on the interface.
VOSS-2128	EAP Security and Authentication EDM tabs display additional information with internal values populated, which is not useful for the end user.	There is no functional impact. Ignore the additional information in EDM. Use the CLI command show eapol port interface to see port status.
VOSS-2207	You cannot configure an SMTP server hostname that begins with a digit. The system displays the following error: Error: Invalid IP Address or Hostname for SMTP server	None.

Issue number	Description	Workaround
VOSS-2208	While performing CFM Layer 2 traceroute between two BEBs using a transit BCB, the transit BCB hop is not seen, if the transit BCB has ISIS adjacencies over FE I3core with both source BEB and destination BEB.	None.
VOSS-2253	Trace level command does not list module IDs when '?' is used.	To get the list of all module IDs, type trace level , and then press Enter .
VOSS-2285	When on BEB, continuously pinging IPv6 neighbor address using CLI command ping -s , ping packets do not drop, but instead return no answer messages.	Restart the ping. Avoid intensive CPU processing.
VOSS-2333	Layer 2 ping to Virtual BMAC (VBMAC) fails, if the VBMAC is reachable using Layer 2 core.	None.
VOSS-2418	When you configure and enable the SLA Mon agent, the SLA Mon server is able to discover it but the agent registration on the switch does not occur.	None.
VOSS-2422	When a BGP Neighbor times out, the following error message occurs: CP1 [03/11/16 13:43:39.084:EST] 0x000b45f2 00000000 GlobalRouter SW ERROR ip_rtdeleteVrf: orec is NULL!	There is no functional impact. Ignore the error message.
VOSS-25476	DvR host entries are visible on DvR Controllers after you issue the clear dvr host-entries command or disable all DvR Controllers within the domain.	Choose one of the following workarounds: <ul style="list-style-type: none"> • Disable and reenabale DvR. • Disable and reenabale IS-IS. • Reenabale DvR Controllers within the domain.
VOSS-2859	You cannot modify the port membership on a protocol-based VLAN using EDM, after it has been created.	Use CLI to provision the port membership on the protocol-based VLAN or delete the protocol-based VLAN, and then re-create it with the correct port member setting.
VOSS-3393	When the SLA Mon agent IP is created on a CLIP interface, the switch provides the CLIP-id as the agent MAC.	There is no functional impact. Use different CLIP IDs to differentiate the SLA Mon agents from the SLA Mon server.
VOSS-4255	If you run IP traceroute from one end host to another end host with a DvR Leaf in between, an intermediate hop will appear as not responding because the Leaf does not have an IP interface to respond. The IP traceroute to the end host will still work.	None.

Issue number	Description	Workaround
VOSS-4728	If you remove and recreate an IS-IS instance on an NNI port with auto-negotiation enabled in addition to vST and R/SMLT enabled, it is possible that the NNI port will briefly become operationally down but does recover quickly. This operational change can lead to a brief traffic loss and possible reconvergence if non-ISIS protocols like OSPF or BGP are also on the NNI port.	If you need to remove and recreate an IS-IS instance on an auto-negotiation enabled NNI port that also has non-ISIS traffic, do so during a maintenance window to minimize possible impact to other non-ISIS traffic.
VOSS-4840	If you run the show fulltech command in an SSH session, do not disable SSH on the system. Doing so can block the SSH session.	None.
VOSS-4912	The VSP 4450 Series does not advertise an LLDP Management TLV.	None.
VOSS-5130	Disabling and immediately enabling IS-IS results in the following log message: <pre>PLSBFIB ERROR: /vob/cb/nd_protocols/plsb/lib/plsbFib.cpp(line 1558)unregisterLocalInfo() local entry does not exist.key(0xfda010000fffa40)</pre>	There is no functional impact. Ignore the error message.
VOSS-5159 & VOSS-5160	If you use a CLIP address as the management IP address, the switch sends out 127.1.0.1 as the source IP address in both SMTP packets and TACACS+ packets.	None.
VOSS-5173	A device on a DvR VLAN cannot authenticate using RADIUS if the RADIUS server is on a DvR VLAN on a DvR Leaf using an in-band management IP address.	Place the RADIUS server in a non-DvR VLAN off a DvR Leaf or DvR Controller.
VOSS-5331	When you enable FHS ND inspection on a VLAN, and an IPv6 interface exists on the same VLAN, the IPv6 host client does not receive a ping response from the VLAN.	None.
VOSS-5603	In a scaled DvR environment (scaled DvR VLANs), you could see a higher CPU utilization while deleting a DvR leaf node from the DvR domain (no dvr leaf). The CPU utilization stays higher for several minutes on that node only and then returns to normal after deleting all the internal VLANs on the leaf node.	It is recommended to use a maintenance window when removing leaf(s) from a DvR domain.

Issue number	Description	Workaround
VOSS-5627	The system does not currently restrict the number of VLANs on which you can simultaneously configure NLB and Directed Broadcast, resulting in resource hogging.	Ensure that you configure NLB and Directed Broadcast on not more than 100 VLANs simultaneously, assuming one NLB cluster for each VLAN. Also, ensure that you configure NLB on a VLAN first, and then Directed Broadcast, so as to not exhaust the NLB and Directed Broadcast shared resources. The shared resources are NLB interfaces and VLANs with Directed Broadcast enabled. The permissible limit for the shared resources is 200.
VOSS-6189	When you connect to EDM using HTTPS in Microsoft Edge or Mozilla Firefox, the configured values for the RADIUS KeepAliveTimer and CFM SBM Mepld do not appear.	Use Internet Explorer when using an HTTPS connection.
VOSS-6822	If the IPsec/IKE software used in the Radius server side is strongSwan, there is a compatibility issue between the network operating system (NOS) and strongSwan in terms of IPv6 Digicert (IKEv1/v2) authentication.	None.
VOSS-6928	On VSP 8000 Series platforms, IPv4 Filters with redirect next hop action do not forward when a default route is not present or a VLAN common to ingress VLAN of the filtered packet is not present.	Configure a default route if possible.
VOSS-7139	DHCPv6 Snooping is not working in an SPB network as the DHCPv6 Snooping entries are not being displayed.	Administrator should add manual entries.
VOSS-7457	The switch can experience an intermittent traffic loss after you disable a Fabric Extend tunnel.	Bounce the tunnel between the devices.
VOSS-7472	EDM shows incorrect guidance for ACL TCP flag mask. EDM reports 0...63 as hexadecimal. CLI correctly shows <0-0x3F 0-63> Mask value <Hex Decimal>. This is a display issue only with no functional impact.	Use CLI to see the correct unit values.
VOSS-7495	The VSP 4450 Series CLI Help text shows an incorrect port for boot config flags linerate-directed-broadcast . The Help text shows 1/48. The correct port is 1/46.	None
VOSS-8424	A fragmented ping from an external device to a switch when the VLAN IP interface is tied to a non-default VRF fails.	None.

Issue number	Description	Workaround
VOSS-8516	Secure Copy (SCP) cannot use 2048-bit public DSA keys from Windows.	Use 1024/2048-bit RSA keys or 1024-bit DSA keys.
VOSS-9516	When you connect to EDM using HTTPS, you can see multiple <code>SSL negotiation with client successful</code> messages during your EDM session. The system displays this message, each time a successful <code>SSL_Handshake</code> occurs between the web browser and the web server. The log file cannot show as many messages as the console and the timing between messages can be different because logging does not occur in real time.	None.
VOSS-9621	On these products, 1G Copper Pluggable auto-negotiation is always enabled after a reboot, despite configuration settings.	If you do not want to use auto-negotiation, disable it after the reboot.
VOSS-9921	Bootup redirection timeout is longer than the UNI port (SMLT) unlock timer. If both vIST nodes boot together in factory default configuration fabric mode or without a nickname, the vIST ports will not enable for up to 4 minutes. During the delay the nickname server is unreachable and vIST is not online.	None.
VOSS-10380	If you enable and configure IPv6 Source Guard and EAPoL on a port, and create and configure a Guest VLAN on the same port without DHCP Snooping and ND-inspection, no error is shown. The port is not added to the Guest VLAN.	Configure DHCP Snooping and ND-inspection are not configured on the Guest VLAN.
VOSS-10381	If you enable and configure IPv6 Source Guard and EAPoL MHSA on a port, and create and configure RAVs for Non-EAP clients on the same port without DHCP Snooping and ND-inspection, no error is shown. The client displays as authenticated into RAV, even when port is not a member of RAV.	None.
VOSS-10412	Removal of the QSFP+ to SFP+ adapter with a 10G pluggable is not detected on the VSP 8404 and VSP 8404C when in non channelized mode.	The QSFP+ to SFP+ adapter and detection works only on ports with channelization enabled.
VOSS-10574	IS-IS sys-name output is not truncated for <code>show isis spbm nick-name</code> or <code>show ip route</code> commands. If a long character sys-name is in use, the full sys-name display can cause misalignment of the output columns.	None.

Issue number	Description	Workaround
VOSS-10815	<p>DvR over SMLT: Traffic is lost at failover on SMLT towards ExtremeXOS or Switch Engine switches. DvR hosts are directly connected to the DvR controllers vIST pair on SMLT LAG and switched-UNIs are dynamically added using Fabric Attach. Only occurs when the access SMLT is LACP MLT and all the ports in the MLT are down.</p> <p>When all ports in the MLT down and an ARP request is received over an NNI link, there is no physical port that can be associated with the ARP request. The ARP entry is learned against NNI link, and MAC syncs from vIST peer or from a non-vIST peer when bouncing vIST.</p>	None.
VOSS-10891	DvR leaf vIST: Wrong rarSmltCheckSmltPeerMac MLT warning displays when the peer vIST MAC address is learned from local	None. rarSmltCheckSmltPeerMac MLT warning has no functional impact. You can ignore the error message.
VOSS-11895	In a vIST SMLT environment where streams are both local and remote, if source and receiver port links are removed and reinserted several times, eventually traffic will not be forwarded to local single-homed receivers on one peer if the traffic is ingressing from the vIST peer over the NNI link. If the stream ingresses locally, it is received by the local UNI receivers.	Disable and re-enable Fabric Multicast (spbm <1-100> multicast enable) on the source VLAN to be able to delete the streams and come back in properly.
VOSS-11943	This release does not support per-port configuration of Application Telemetry. Because the feature is enabled globally and VSP 7432CQ supports 32 100 Gbps ports, an undesirable condition could be encountered when an exceeded amount of Application Telemetry mirrored packets are sent to the collector.	None.
VOSS-12330	When accessing the on-switch RESTCONF API documentation in a web browser, the page does not render correctly.	Ensure you include the trailing slash (/) in the URL: <code>http(s)://<ip-address>:8080/apps/restconfdoc/</code> . For more information, see Fabric Engine User Guide .

Issue number	Description	Workaround
VOSS-12405	To reach a VM, all front panel traffic must travel through an Insight port, which is a 10 Gbps port. If front panel port traffic is over 10 Gbps, this situation represents an over subscription on the Insight port and some of the packets will be dropped. As a result, ExtremeCloud IQ - Site Engine can lose connectivity to the Analytics engine if Application Telemetry is enabled.	None.
VOSS-13159	The ixgbev Ethernet device driver within the TPVM does not correctly handle the interface MTU setting. Specifically, if you configure the interface in SR-IOV mode, packets larger than the MTU size are allowed.	To avoid this problem, configure the desired MTU size on both the relevant front-panel port and Insight port from the NOS CLI.
VOSS-13463	Out port statistics for MLT port interfaces are not accurate.	Use the command <code>show io nic-counters</code> to display detailed port stats and error info on XA1400 Series.
VOSS-13667	An intermittent issue in SMLT environments, where ARPs or IPv6 neighbors are resolved with delay can cause a transient traffic loss for the affected IPv6 neighbors. The situation auto-corrects.	None.
VOSS-13680	Interface error statistics display is inaccurate in certain scenarios.	Use the command <code>show io nic-counters</code> to display detailed port stats and error info on XA1400 Series.
VOSS-13681	QoS: show qos cosq-stats cpu-port command output is not supported.	Use the command <code>show io cpu-cosq-counters</code> to display detailed cosq-stats on XA1400 Series.
VOSS-13693	QoS: Traffic can egress out of the queue at a different ratio than the default configuration. After the guaranteed traffic rate is served to all egress port queues, any excess bandwidth is shared equally to all queues instead of distributing on weight assigned to each queue.	None.
VOSS-13717 VOSS-14393 VOSS-14972	Link on remote side doesn't go down after admin shut on XA1400 while using 10G DAC or a 4x10 - 40 G breakout DAC. On the XA1400 side link goes down but Link LED shows as up. Both 10G and 4x10G DAC are not fully supported because of this issue	None for DAC and breakout cables. Because of this issue, the following optical transceivers are not supported: <ul style="list-style-type: none"> • AA1404036-E6 • AA1404042-E6 • C9799X4-5M
VOSS-13794	You cannot use SFTP to transfer files larger than 2 GB to the switch.	Use SCP.

Issue number	Description	Workaround
VOSS-13904 VOSS-13932 VOSS-16503	VSP 4900 Series has 2 GB memory in a 64-bit system so the RESTCONF VLAN scaling number is smaller than on VSP 7400 Series, which has 16 GB physical memory. Using RESTCONF on VSP4900-48P or VSP4900-24S reduces the number of port-based VLANs on those platforms: <ul style="list-style-type: none"> • 2,000 for VSP4900-48P with RESTCONF • 1,000 for VSP4900-24S with RESTCONF 	None.
VOSS-13947	After you enable MSTP-Fabric Connect Multi Homing (spbm 1 stp-multi-homing enable), you cannot view the configuration, role, or statistics for the STP virtual port.	None.
VOSS-13974	When an 8408QQ ESM has more than two channelized ports and is rebooted, the MKA MACsec sessions on the other cards in the same box could toggle. This issue is not seen if one or two ports are channelized on the same card.	None.
VOSS-14150	CLI remote console might stop wrapping text after some usage.	Reset the CLI window or open a new remote console window.
VOSS-14391	On an VSP 8404C switch using an 8424XT ESM, on a port with MACsec connectivity, if you set Auto-Negotiation advertisements to 1000-full, and then subsequently set the advertisement to 10000-full, the link will not come up.	To avoid this issue, set the Auto-Negotiation advertisements directly to 10000-full. If you have experienced the issue, shut the port down and bring it back up.
VOSS-14494	Layer 2 VSN and Layer 3 VSN UNI to NNI traffic between two Backbone Edge Bridges does not hash to different ports of a MLT network-to-network interface. MLT hashing for XA1400 devices occurs after the mac-in-mac encapsulation is done. The hash keys used are the Backbone destination and Backbone source MAC addresses (BMAC DA and BMAC SA) in the Mac-in-Mac header. Even for the Transit BCB case on XA 1400 devices for NNI to NNI traffic, the MLT hash keys used are the Backbone destination and Backbone source MAC addresses (BMAC DA and BMAC SA) in the Mac-in-Mac header.	None.

Issue number	Description	Workaround
VOSS-14515	<p>Console output errors and warnings are shown during an XA1400 Series reboot, such as:</p> <ul style="list-style-type: none"> • error: no such device: ((hd0,gpt1)/EFI/BOOT)/EFI/BOOT/grub.cfg. error: file `/EFI/BOOT/grubenv' not found • error: no suitable video mode found. • [0.727012] ACPI: No IRQ available for PCI Interrupt Link [LNKS]. Try pci=noacpi or acpi=off • exportfs: can't open /etc/exports for reading • KCORE: WARNING can't find /boot/b/ulmage-gemini.bin. No kexec kernel will be configured. 	None. The errors or warnings are host OS or guest OS related with no functional impact and can be ignored.
VOSS-14597	Ping (originated from local CP) fails for jumbo frames on Layer 3 VSN interface.	None.
VOSS-14616	<p>Seeing Queue buffer usage logs when changing the logical interface source IP with 64 tunnels.</p> <p>When changing the source IP with 64 tunnels, seeing "GlobalRouter CPU INFO CPP: 60 percent of fbufs are in use: 0 in Tx queue,1843 in RxQueue0 0 in RxQueue1 0 in RxQueue2 0 in RxQueue3 0 in RxQueue4 0 in RxQueue5 0 in RxQueue6 0 in RxQueue7 ".</p>	None.
VOSS-14805 VOSS-15305	<p>The following transceivers are not supported on XA1400 Series switches:</p> <ul style="list-style-type: none"> • 10 Gb Bidirectional 40 km SFP+ Module (10GB-BX40-D and 10GBBX40-U) • 1000BASE-BX10 Bidirectional 10 km DDI SFP Modules (AA1419069-E6 and AA1419070-E6) 	Use only supported transceivers.
VOSS-15079	The Extreme Networks 10 meter SFP+ passive copper DAC (Model Number 10307) does not function on ports 2/3 and 2/4 of the VIM5-4X.	Use the Extreme Networks SFP+ active optical DAC (Model Number AA1403018-E6) with the VIM5-4X.
VOSS-15112	BFD sessions associated with static routes could flap one time before remaining up, when shutting down and bringing back up a BFD peer port.	None. Ignore the extra BFD session flap.

Issue number	Description	Workaround
VOSS-15313	On a VSP 8404C switch using an 8424XT ESM, on a link with MACsec connectivity on both ends, and Auto-Negotiation advertisements set to 10000-full, the link will not come back up if the ESM is hot-swapped or the slot is reset.	To avoid this issue, disable MACsec prior to the hot swap or reset, and then re-enable. If you have experienced the issue, shut either one of the link ports down and bring it back up.
VOSS-15391	An SNMP walk on the rcIgmpSnoopTraceTable table will fail with an OID not increasing error. CLI and EDM are unaffected by this issue.	None.
VOSS-15463	XA1440 and XA1480 switches can experience intermittent Link Up and Link Down transitions on the 10/100/1000BASE-T Ethernet ports upon booting.	No workaround, but there is no functional impact.
VOSS-15541	You can experience temporary traffic loss when shutting down an LACP SMLT port (and therefore causing the local SMLT to go down), in a network with scaled Multicast traffic over an SPB cloud, while the datapath processes all dpm letter messages during LCAP recovery. This slow LACP recovery situation is only seen with scaled Multicast traffic over an SPB cloud.	Use static MLTs.
VOSS-15812	Layer 3 VSN IPv4 BGP (and static) routes having their next-hops resolved using IS-IS routes could result in traffic loss.	Choose the following workarounds, based on your deployment and needs: <ul style="list-style-type: none"> • Use static routes to reach the loopbacks used as BGP peers, (static routes having better preference than IS-IS); use static routes with next-hops reachable on the UNI side (L2VSN). • Use OSPF to reach the loopbacks used as BGP peers, but take care to ensure that the OSPF route towards the BGP peer is chosen as the “best route” (as IS-IS has a better preference than OSPF). There are several ways to accomplish this— either don’t redistribute that route in IS-IS if it is not needed, or control the redistribution with a route-map, etc. • Have BGP peers reachable directly using a C-VLAN; do not use loopback interfaces as BGP peer addresses. • If none of the above workaround scenarios are suitable for your deployment, do not use internal Border Gateway Protocol (iBGP) peering.

Issue number	Description	Workaround
VOSS-15878	VSP 4900 Series, VSP 7400 Series do not boot with just the serial console cable connected and no terminating device, for example, a terminal server, PC, or Mac.	Either attach terminal equipment or disconnect the console cable.
VOSS-16221	Layer 2 ping does not work for packets larger than 1300 on an XA1400 Series.	Use Layer 2 ping with packets smaller than 1300 bytes.
VOSS-16365	Running the command show pluggable-optical-module detail on an XA1400 Series device is highly CPU intensive to read and reply with the EEPROM details. Due to a delay in ethtool response, a watchdog miss event can occur and the event is recorded in the <code>/intflash/wd_stats/1/wd_stats.ssio.1.log</code> file. This scenario occurs more often if 10Gb SFP+ optics with DDM capability are installed.	None. The high CPU usage and response delay for this command is expected and cannot be resolved. No console log is generated. When the scenario occurs, the Watchdog outage is approximately 5 seconds.
VOSS-16436	Using the console connection on an XA1400 Series device while running a show command with large data output can result in drops of processing control packets.	Use Telnet or SSH connectivity instead of console connection.
VOSS-16951	On a VSP4900-48P, VSP4900-24S and VSP 7400 Series devices, if you run the show boot config sio CLI command before you have configured the baud rate, the output of the command is empty.	Configure the baud rate before you run the show boot config sio command. The only supported baud rate for these devices is 115200.
VOSS-16971	On VSP4900-24S, VSP4900-24XE, and VSP4900-12MXU-12XE devices, and on the VIM5-4XE, if a copper SFP is plugged in with the cable inserted and the remote end is also plugged in, the peer box could see a link flap and take 6-8 seconds to link up.	First, plug in the SFP, and then insert the cable. The link up then happens in 3-4 seconds.
VOSS-17002	For ingress packets that are larger than the system MTU size on XA1400 Series ports 1/1 through 1/4, error counters do not increment in the show interfaces gigabitethernet error CLI command.	Use the show io nic-counters CLI command to verify if the <code>tx_error</code> counters are getting incremented. If they are getting incremented, the packets are getting dropped at egress. If they are not getting incremented, the packets are getting forwarded.
VOSS-17523	If an FE tunnel goes down between two connected XA1400 Series devices, an MTU Warning console message is logged if a ping request is issued while the tunnel is down.	You can safely ignore this warning message.

Issue number	Description	Workaround
VOSS-17567	Do not use the inter-vrf /32 static routes defined with a next-hop IP address that resides in a different destination next-hop-vrf context.	None.
VOSS-18023	<p>The management port on the 5520 switch does not support Auto-MDIX (the automatic detection of transmit and received twisted pairs). As a best practice, enable the default auto-negotiation setting on the management port. Because the management port does not support Auto-MDIX, when auto-negotiation is disabled, a crossover cable might be necessary to have the port link up and pass traffic.</p> <p>Note: If the peer device supports Auto-MDIX, then either a straight through or crossover will work. The issue occurs only if both ends of the connection do not support Auto-MDIX.</p>	None.
VOSS-18238	When a management VLAN with DHCP is used to reach a RADIUS server, and the RADIUS server cannot be reached, the system waits for 15 minutes before attempting to reach the RADIUS server again. This is true even if the RADIUS server becomes reachable before the 15 minutes have elapsed.	None.
VOSS-18278	<p>On the 5520 switch, when you make any change relating to port speed, the port statistics are cleared. This applies to all front panel fiber and copper ports as well as VIM ports. The following are examples of changes relating to port speed:</p> <ul style="list-style-type: none"> • Changing the auto-negotiation configuration settings on a copper port • Different negotiated speed on a copper port • Changing out an optical device for one having a different speed, for example changing from 1 Gb to 10 Gb 	None.

Issue number	Description	Workaround
VOSS-18360	This is an intermittent issue on the VSP 7400 Series with no impact to functionality, ISIS is disabled while the show fulltech command is running on a telnet session. Due to this the fulltech command will not find the expected I-SID value, as it is removed by the no isis command.	None.
VOSS-19212	After upgrading a VSP 7432CQ switch to VOSS 8.2.5 and rebooting, the presence of a faulty power supply unit will cause the system to terminate. A message in the debug log will report that the software could not read the contents of the power supply's EEPROM (<i>carbonatelib_ps_read_eeprom</i> operation).	Replace the power supply unit in the switch.
VOSS-19260	Port mirroring does not work on port 1/s1 of VSP 7400-48Y if the connection type is OVS/SR-IOV.	Use a connection type of VT-d for port 1/s1.
VOSS-19827	LLDP IPv6 neighbors do not display in EDM. LLDP IPv6 is only supported in CLI.	To display LLDP IPv6 neighbors, use the show lldp neighbor summary command.
VOSS-20115	You cannot change the management VLAN interface discovered on XA1400 Series in ExtremeCloud IQ - Site Engine as part of Zero Touch Provisioning Plus (ZTP+). XA1400 Series does not support the OOB interface. You can only use the discovered interface and change other configuration values.	On XA1400 Series, use the discovered interface within ExtremeCloud IQ - Site Engine for basic onboarding. Use either ExtremeCloud IQ - Site Engine or CLI to complete the remaining configuration.
VOSS-20200	For VSP 8404C, if you remove and insert an Ethernet Switch Module (ESM), which has NNI ports that are members in an LACP-dynamic MLT, some ports are intermittently missing in the dynamic MLT after the ESM insertion. Traffic is affected for streams that need to exit the NNI links over the dynamic MLT for the missing ports. Rebooting the switch returns the ports to the dynamic MLT.	None.
VOSS-20227	On XA1400 Series, the VOSS OS time does not synchronize to the real time clock (RTC) after system reboot. After the switch completely boots, NTP synchronization occurs and the VOSS OS has the correct time. The OS time can be incorrect for up to two minutes after system reboot.	None.

Issue number	Description	Workaround
VOSS-20455	<p>As the switch starts, it can display the following log messages due to incomplete initialization of the management stack when trying to send the first RADIUS packet:</p> <ul style="list-style-type: none"> • 1 2021-02-17T23:32:16.810+01:00 DIST-H9-E3.1-01 CP1 - 0x000a45ae - 00000000 GlobalRouter RADIUS ERROR rad_sendRequest: unable to send a UDP packet. error 51, S_errno_ENETUNREACH • 1 2021-02-17T23:32:16.811+01:00 DIST-H9-E3.1-01 CP1 - 0x000a45ac - 00000000 GlobalRouter RADIUS ERROR rad_processPendingRequest: unable to send request 	None. This issue has no functional impact.
VOSS-20456	Although the Management Router is not supported in the NOS, you can add a static route for VRF 512 using EDM. The route does not become active even if the next-hop address is reachable from the OOB management interface.	None. This issue has no functional impact.
VOSS-21097	In Multi-Area where vIST peers are boundary nodes, vIST can briefly flap during connection formation when IS-IS is disabled and then reenabled on both vIST peers.	None.
VOSS-21123	Routers on UNIs of VSP 7400 vIST peers cannot ping each other.	Add a static ARP for the Router of the VIST peer.
VOSS-21233	Clearing DvR host entries in a highly scaled Multi-Area DvR environment can trigger DBSYNC WARNING messages (0x00390606 - 00000000 GlobalRouter DBSYNC WARNING Message queue length from DB Sync to tMain reached warning threshold) but these can be expected in a scaled environment and are not a malfunction.	None.
VOSS-21964	When using Windows SCP application on a switch to transfer a file, an error message displays even if a file transfers successfully.	
VOSS-22255	Ping, which originates from a local CP, fails for ICMP packets bigger than 1500 sent from Layer 3 VSN interface.	Initiate ping with packets size smaller than 1500.

Issue number	Description	Workaround
VOSS-22522	RESTCONF is delayed in a scaled setup with 2,000 VLANs.	None.
VOSS-22858	LLDP neighbor should not be discovered with mismatch in MKA MACsec on 5520 Series ports.	Disable MKA on both sides or shut down the port on both sides.
VOSS-23146	Multi-area DvR/SPBM configuration: Timeout: No response message is returned during snmpwalk on one of the DvR controllers.	Run the snmpwalk command with an increased timeout. You can also run snmpwalk for a specific object.
VOSS-23181	When you enable the boot config flags macsec command, the indiscard counter increments on SPBM-enabled ports.	None. There is no functional impact.
VOSS-23216	If you do not enable the DvR interface when you configure a dvr-one-ip interface, the dvr-one-ip interface does not display when you issue the show dvr interfaces command.	Enable the DvR interface.
VOSS-23229	In an E-Tree scenario, IPv6 packets are forwarded between isolated ports on 5520 Series, 5420 Series, and VSP 7400 Series.	None.
VOSS-24777	In the following port configurations on 5520 Series, 5420 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, and VSP 8400 Series inVSN ACL entries match ingressing packets that have the same VID as the VLAN associated with the ACL I-SID even if the ACL inVSN I-SID is different: <ul style="list-style-type: none"> on an S-UNI port without a platform VLAN on a T-UNI port VLAN 	None.
VOSS-24872	If the collector reachability path changes for Application Telemetry, it is not reflected properly in CLI. Packets remain mirrored towards the correct path but CLI does not reflect the next hop.	None. There is no functional impact.
VOSS-25078	MAC addresses learned on a Switched UNI (S-UNI) port cannot be flushed.	None.
VOSS-25023	5520 Series, 5420 Series, and 5320 Series platforms can reach 100% CPU utilization during inband transfer (FTP, SFTP, and SCP).	None.

Issue number	Description	Workaround
VOSS-25162	RESTCONF ARP and MAC data: on 5x20 switches with 5K ARP entries and 5K MAC entries, it takes approximately 1 minute to retrieve data. The time increases based on the number of entries. The same occurs on VSP 7400 Series with over 15K entries.	None.
VOSS-25225	On 5320 Series, the four highest SFP+ ports are available at 10 Gbps with Trial Licenses. After license expiration, the port speeds drop to 1 Gbps.	Use the extend-time-period command prior to the expiration of the Trial License.
VOSS-25288	Secure boot information for 5720 Series does not display when you issue the show sys-info command.	None.
VOSS-25728	You cannot assign a second disk to the second virtual service on the following switches: <ul style="list-style-type: none"> • VSP 4900 Series • VSP 7400 Series • 5720 Series 	None.
VOSS-25874	Intermittent issue seen on CFIT rack that causes inconsistency in show output.	None.
VOSS-25959	On the VSP 4900 Series, VSP 7400 Series, and 5720 Series, the virtual service does not operate properly when you configure <i>e1000</i> Network Interface Card (NIC) type for SR-IOV and VT-d connect types.	None.
VOSS-26028	On the VSP 4900 Series, VSP 7400 Series, and 5720 Series, the virtual service does not operate properly when you configure more than 16 virtual ports per Extreme Integrated Application Hosting port.	None.
VOSS-26032	NNI port remains in STP blocking state in a very specific scenario and configuration.	Bounce the NNI port.
VOSS-26092	On the VSP 8400 Series, MKA does not operate after you issue the slot reset command.	As a workaround, issue the reset command to reset your switch.
VOSS-26099	MACsec Key Agreement (MKA) MACsec does not operate properly when you enable and disable MKA MACsec on the port 15-20 times.	None.
VOSS-26122	Intermittently, some CLI commands related to sFlow functionality do not display in the CLI log.	None.

Issue number	Description	Workaround
VOSS-26134	On the VSP 7200 Series, ports link flap one time when the switch boots and after you issue the shutdown command.	None.
VOSS-26151	MACsec Key Agreement (MKA) does not operate between Fabric Engine 5520 Series and 5720 Series switches and ExtremeXOS 5520 Series and 5720 Series switches when you use GCM-AES-256 MACsec encryption cipher suite on copper ports.	As a workaround, use GCM-AES-128 MACsec encryption cipher suite to connect Fabric Engine 5520 Series and 5720 Series switches and Switch Engine 5520 Series and 5720 Series switches.
VOSS-26526	After you format a USB drive and issue the ls command, the current date and time does not display.	None.
VOSS-26527	Intermittently, the show sys-info command does not display the correct part number or serial number for the 2000 W AC PoE power supply (Model XN-ACPWR-2000W with front-to-back ventilation airflow).	None.
VOSS-26665	Password hash sha2 is present in show running-config and save config . This is the default value.	None.
VOSS-26692	The entry for VLAN used to send/receive VXLAN packets to/from FIGW (for IPsec encapsulation) is missing from my_station_tcam table. In this case, traffic over the corresponding FE tunnel is lost.	Shut/no shut of the used sideband port fixes the problem.
VOSS-26822	Configuration tab for Ports 53-54 (VSP 7400-48Y) cannot be accessed from the first attempt.	Select menu options on your Mozilla Firefox browser. Alternatively, use another browser: Google Chrome, Safari, or Microsoft Edge.
VOSS-26831	Device not able to complete trap registration with ExtremeCloud IQ - Site Engine when onboarding with ZTP+.	Use the default Trap profile when using Trap registration with auto onboarding in ExtremeCloud IQ - Site Engine.
VOSS-26884	AP is assigned to an Unregistered rule instead of Wifi Mgmt on a 22.9 version NAC.	None.
VOSS-27235	If you delete a VLAN IP interface, the switch does not delete the associated DvR gateway IP address.	Manually delete the DvR gateway IP address.

Restrictions and Expected Behaviors

This section lists known restrictions and expected behaviors that can first appear to be issues.

For Port Mirroring considerations and restrictions, see [Fabric Engine User Guide](#).

General Restrictions and Expected Behaviors

The following table provides a description of the restriction or behavior.

Table 30: General restrictions

Issue number	Description	Workaround
—	If you access the Extreme Integrated Application Hosting virtual machine using virtual-service tpvm console and use the Nano text editor inside the console access, the command ^o<cr> does not write the file to disk.	None.
VOSS-7	Even when you change the LLDP mode of an interface from CDP to LLDP, if the remote side sends CDP packets, the switch accepts them and refreshes the existing CDP neighbor entry.	Disable LLDP on the interface first, and then disable CDP and re-enable LLDP.
VOSS-687	EDM and CLI show different local preference values for a BGP IPv6 route. EDM displays path attributes as received and stored in the BGP subsystem. If the attribute is from an eBGP peer, the local preference displays as zero. CLI displays path attributes associated with the route entry, which can be modified by a policy. If a route policy is not configured, the local preference shows the default value of 100.	None.
VOSS-1954	After you log in to EDM, if you try to refresh the page by clicking on the refresh button in the browser toolbar, it will redirect to a blank page. This issue happens only for the very first attempt and only in Firefox.	To refresh the page and avoid this issue, use the EDM refresh button instead of the browser refresh button. If you do encounter this issue, place your cursor in the address bar of the browser, and press Enter . This will return you to the EDM home page.
VOSS-2166	The IPsec security association (SA) configuration has a NULL Encryption option under the Encrpt-algo parameter. Currently, you must fill the encrptKey and keyLength sub-parameters to set this option; however, these values are not used for actual IPsec processing as it is a NULL encryption option. The NULL option is required to interoperate with other vendors whose IPsec solution only supports that mode for encryption.	There is no functional impact due to this configuration and it only leads to an unnecessary configuration step. No workaround required.
VOSS-21946	When you create a vrf using the POSTMAN API platform, special characters, such as \\ \\ \\ and ### included in the URL are ignored.	None.

Table 30: General restrictions (continued)

Issue number	Description	Workaround
VOSS-2185	MAC move of the client to the new port does not automatically happen when you move a Non-EAP client authenticated on a specific port to another EAPoL or Non-EAP enabled port.	As a workaround, perform one of the following tasks: <ul style="list-style-type: none"> • Clear the non-EAP session on the port that the client is first authenticated on, before you move the client to another port. • Create a VLAN on the switch with the same VLAN ID as that dynamically assigned by the RADIUS server during client authentication. Use the command vlan create <2-4059> type port-mstprstp <0-63>. Ensure that the new port is a member of this VLAN.
VOSS-5197	A BGP peer-group is uniquely identified by its name and not by its index. It is possible that the index that is configured for a peer-group changes between system reboots; however this has no functional impact.	None.
VOSS-7553	Option to configure the default queue profile rate-limit and weight values are inconsistent between EDM and CLI. Option to configure default values is missing in EDM.	None.
VOSS-7640	The same route is learned via multiple IPv6 routing protocols (a combination of two of the following : RIPng, OSPFv3 and BGPv6). In this specific case, an eBGP (current best - preference 45) route is replaced by and iBGP (preference 175) which in turn is replaced by and OSPFv3 (external 2) route (preference 125).	None.
VOSS-7647	With peer group configuration, you cannot configure Update Source interface with IPv6 loopback address in EDM.	Use CLI.
VOSS-9174	OVSDB remote VTEP and MAC details can take between 5 to 10 minutes to populate and display after a HW-VTEP reboots.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.
VOSS-9462	OVSDB VNID I-SID MAC bindings are not populated on HW-VTEPs after configuration changes.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.

Table 30: General restrictions (continued)

Issue number	Description	Workaround
VOSS-10168	The system CLI does not prevent you from using the same IP address for the VXLAN Gateway hardware VTEP replication remote peer IP and OOB Management IP.	Manually check the IP configured as the OOB Management IP. Do not use the OOB Management IP address as the replication remote peer IP address.
VOSS-11817	The OVS connect-type for virtual service Vports is designed in such a way that it connects to any generic virtual machine (VM) guest OS version using readily available Ethernet device drivers. This design approach provides initial connectivity to the VM in a consistent manner. A consequence of this approach is that Vports created with connect-type OVS will show up as 1 Gbps interfaces in the VM even though the underlying Ethernet connection supports 10 Gbps .	If additional performance is desired, upgrade the VM guest OS with an Ethernet device driver that supports 10 Gbps interfaces.
VOSS-12151	If logical switch has only hardware ports binding, and not VM behind software VTEP, Broadcast, Unknown Unicast, and Multicast (BUM) traffic does not flow between host behind two hardware VTEP. The NSX replicator node handles the BUM traffic. NSX does not create the replicator node unless a VM is present. In an OVSDB topology, it is expected that at least one VM connects to the software VTEP. This issue is an NSX-imposed limitation.	After you connect the VM to the software VTEP, the issue is not seen.
VOSS-12395	You cannot use the following cables on 10 Gb fiber interfaces, or 40 Gb channelized interfaces, with the QSA28 adapter: <ul style="list-style-type: none"> • 1, 3, and 5 meter QSFP28 25 Gb DAC • 20 meter QSFP28 25 Gb AOC 	n/a
VOSS-17871	Starting with VOSS 8.1.5, internal system updates have resulted in a more accurate accounting of memory utilization. This can result in a higher baseline memory utilization reported although actual memory usage is not impacted.	Update any network management alarms that are triggered by value with the new baseline.
VOSS-18523	When you configure a port using Zero Touch Provisioning Plus (ZTP+) with ExtremeCloud IQ - Site Engine, the port cannot be part of both a tagged VLAN and an untagged VLAN.	n/a

Table 30: General restrictions (continued)

Issue number	Description	Workaround
VOSS-18409	On the XA1400 Series switches, only one Central Processing Unit (CPU) core is assigned for control plane protocol processing. In a highly scaled scenario, a port toggling or negative scenario keeps the CPU core busy in updating the software datapath entries. Similarly, some show CLI commands that require a lot of data gathering keep the CPU core busy. In such a scenario, the main task which is responsible for handling protocol packets like Bidirectional Forwarding Detection, Intermediate-System-to-Intermediate-System, Virtual Link Aggregation Control Protocol, and so on is busy.	For scaled scenarios on XA1400 Series switches, the CLI commands that have large sections of output, for example, show fulltech, show io spb tables, and show tech, the output must be redirected into a file.
VOSS-18774	SSL negotiation fails when using OpenSSL client version 1.1.1. With OpenSSL 1.1.1, the server-name extension is used. This extension needs to equal the domain name in the server certificate, otherwise the certificate lookup on the server fails because the FIPS 140-2 certified cryptographic module processes the server-name extension.	Can connect using: <pre>bash# openssl s_client -connect <domain-name>:443</pre>
VOSS-18851	Do not define a static route in which the NextHop definition uses an Inter-VRF redistributed route. Such a definition would require the system to perform a double lookup. When you attempt to define a static route in this way, an error message is generated.	Define the static route in such a way that it does not require Inter-VRF redistributed routing.
VOSS-21620	When interior nodes are running software earlier than Release 8.4 and a Multi-area takeover occurs between the boundary nodes (when the non-designated boundary node transitions to designated) in the network, the interior nodes might detect a false duplicate case between the stale LSP of the old virtual node and the new virtual node. This has no functional impact in the network.	n/a

Table 30: General restrictions (continued)

Issue number	Description	Workaround
wi01068569	The system displays a warning message that routes will not inject until the apply command is issued after the enable command. The warning applies only after you enable redistribution, and not after you disable redistribution. For example: Switch:1 (config)#isis apply redistribute direct vrf 2	n/a
wi01112491	IS-IS enabled ports cannot be added to an MLT. The current release does not support this configuration.	n/a
wi01122478	Stale SNMP server community entries for different VRFs appear after reboot with no VRFs. On a node with a valid configuration file saved with more than the default vrf0, SNMP community entries for that VRF are created and maintained in a separate text file, <code>snmp_comm.txt</code> , on every boot. The node reads this file and updates the SNMP communities available on the node. As a result, if you boot a configuration that has no VRFs, you can still see SNMP community entries for VRFs other than the globalRouter vrf0 .	n/a
wi01137195	A static multicast group cannot be configured on a Layer 2 VLAN before enabling IGMP snooping on the VLAN. After IGMP snooping is enabled on the Layer 2 VLAN for the first time, static multicast group configuration is allowed, even when IGMP snooping is disabled later on that Layer 2 VLAN.	n/a
wi01141638	When a VLAN with 1000 multicast senders is deleted, the console or Telnet session stops responding and SNMP requests time out for up to 2 minutes.	n/a

Table 30: General restrictions (continued)

Issue number	Description	Workaround
wi01142142	When a multicast sender moves from one port to another within the same BEB or from one vIST peer BEB to another, with the old port operationally up, the source port information in the output of the show ip igmp sender command is not updated with new sender port information.	<p>You can perform one of the following workarounds:</p> <ul style="list-style-type: none"> On an IGMP snoop-enabled interface, you can flush IGMP sender records. <p>Caution: Flushing sender records can cause a transient traffic loss.</p> <ul style="list-style-type: none"> On an IGMP-enabled Layer 3 interface, you can toggle the IGMP state. <p>Caution: Expect traffic loss until IGMP records are built after toggling the IGMP state.</p>
wi01145099	IP multicast packets with a time-to-live (TTL) equal to 1 are not switched across the SPB cloud over a Layer 2 VSN. They are dropped by the ingress BEB.	To prevent IP multicast packets from being dropped, configure multicast senders to send traffic with TTL greater than 1.
wi01159075	VSP 4450GTX-HT-PWR+: Mirroring functionality is not working for RSTP BPDUs.	None.
wi01171670	Telnet packets get encrypted on MACsec-enabled ports.	None.
wi01198872	On VSP 4450 Series, a loss of learned MAC addresses occurs in a vIST setup beyond 10k addresses. In a SPB setup the MAC learning is limited to 13k MAC addresses, due to the limitation of the internal architecture when using SPB. Moreover, as vIST uses SPB and due to the way vIST synchronizes MAC addresses with a vIST pair, the MAC learning in a vIST setup is limited to 10K Mac addresses.	None.
wi01210217	The command show eapol auth-stats displays LAST-SRC-MAC for NEAP sessions incorrectly.	n/a
wi01211415	In addition to the fan modules, each power supply also has a fan. The power supply stops working if a power supply fan fails, but there is no LED or software warning that indicates this failure.	Try to recover the power supply fan by resetting the switch. If the fan does not recover, then replace the faulty power supply.

Table 30: General restrictions (continued)

Issue number	Description	Workaround
wi01212034	When you disable EAPoL globally: <ul style="list-style-type: none"> Traffic is allowed for static MAC configured on EAPoL enabled port without authentication. Static MAC config added for authenticated NEAP client is lost. 	n/a
wi01212247	BGP tends to have many routes. Frequent additions or deletions impact network connectivity. To prevent frequent additions or deletions, reflected routes are not withdrawn from client 2 even though they are withdrawn from client 1. Disabling route-reflection can create a black hole in the network.	Bounce the BGP protocol globally.
wi01212585	LED blinking in EDM is representative of, but not identical to, the actual LED blinking rates on the switch.	n/a
wi01213040	When you disable auto-negotiation on both sides, the 10 Gbps copper link does not come up.	n/a
wi01213066 wi01213374	EAP and NEAP are not supported on brouter ports.	n/a
wi01213336	When you configure tx mode port mirroring on T-UNI and SPBM NNI ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports appear on the mirror destination port, although they do not egress the mirror source port. This is because tx mode port mirroring happens on the mirror source port before the source port squelching logic drops the packets at the egress port.	n/a
wi01219658	The command show khi port-statistics does not display the count for NNI ingress control packets going to the CP.	n/a
wi01219295	SPBM QOS: Egress UNI port does not follow port QOS with ingress NNI port and Mac-in-Mac incoming packets.	n/a
wi01223526	ISIS logs duplicate system ID only when the device is a direct neighbor.	n/a
wi01223557	Multicast outage occurs on LACP MLT when simplified vIST peer is rebooted.	You can perform one of the following workarounds: <ul style="list-style-type: none"> Enable PIM on the edge. Ensure that IST peers are either RP or DR but not both.

Table 30: General restrictions (continued)

Issue number	Description	Workaround
wi01224683 wi01224689	Additional link bounce can occur on 10 Gbps ports when toggling links or during cable re-insertion. Additional link bounce can occur with 40 Gbps optical cables and 40 Gbps break-out cables, when toggling links or during cable re-insertion.	n/a
wi01229417	Origination and termination of IPv6 6-in-4 tunnel is not supported on a node with vIST enabled.	None.
wi01232578	When SSH keyboard-interactive-auth mode is enabled, the server generates the password prompt to be displayed and sends it to the SSH client. The server always sends an expanded format of the IPv6 address. When SSH keyboard-interactive-auth mode is disabled and password-auth is enabled, the client itself generates the password prompt, and it displays the IPv6 address format used in the ssh command.	None.
wi01234289	HTTP management of the ONA is not supported when it is deployed with a VSP 4450 Series device.	None.
VOSS-26218	In a scaled environment, running the show io 12-tables command reiteratively can cause the switch to reboot.	For scaled scenarios, do not run the show io 12-tables command in a loop.

Filter Restrictions

The following table identifies known restrictions.

Table 31: ACL restrictions

Applies To	Restriction
All platforms	Only port-based ACLs are supported on egress. VLAN-based ACLs are not supported.
All platforms	IPv6 ingress and IPv6 egress QoS ACL/filters are not supported. Note: IPv6 ACL DSCP Remarking is supported.
All platforms	Control packet action is not supported on InVSN Filter or IPv6 filters generally.
All platforms	IPv4/IPv6 VLAN based ACL filters will be applied on traffic received on all the ports if it matches VLAN ID associated with the ACL.
All platforms	Scaling numbers are reduced for IPv6 filters.
All platforms	The InVSN Filter does supports IP Shortcut traffic only on both UNI and NNI ports, but does not support IP Shortcut traffic on UNI ports only and NNI ports only.

Table 31: ACL restrictions (continued)

Applies To	Restriction
All platforms	The InVSN Filter does not filter packets that arrive on NNI ingress ports but are bridged to other NNI ports or are for transit traffic.
All platforms	You can insert an InVSN ACL type for a Switched UNI only if the Switched UNI I-SID is associated with a platform VLAN.

Table 32: ACE restrictions

Applies To	Restriction
All platforms	When an ACE with action count is disabled, the statistics associated with the ACE are reset.
All platforms	Only security ACEs are supported on egress. QoS ACEs are not supported.
All platforms	ICMP type code qualifier is supported only on ingress filters.
All platforms	For port-based ACLs, you can configure VLAN qualifiers. Configuring port qualifiers are not permitted.
All platforms	For VLAN-based ACLs, you can configure port qualifiers. Configuring VLAN qualifiers are not permitted.
All platforms	Egress QoS filters are not supported for IPv6 filters.
All platforms	Source/Destination MAC addresses cannot be added as attributes for IPv6 filters ACEs.



Resolved Issues this Release

This release incorporates all fixes from prior releases, up to and including the following releases:

- VOSS 8.5.3
- Fabric Engine 8.8.1

Issue number	Description
VOSS-18477	On the VSP 4900 Series, an intermittent traffic loss over the FE tunnels, in SMLT contexts, occurs for a few seconds, when you read ports to the SMLT trunk.
VOSS-24771	When you configure the macsec connectivity-association name to the maximum of 16 characters using CLI, the connectivity association name attached to the port or interface does not display in EDM.
VOSS-25910	In highly scaled networks where the PSNP might be processed with delay, because the PSNPs were not processed yet for the LSPs sent on the already up adjacencies when an additional adjacency comes up, the LSPs, local or forwarded from other nodes, will be resent on all the NNIs, which creates an overhead.
VOSS-26579	When an NNI is created using run spbm , that port is still in VLAN 1; STP is enabled in CIST on that port.
VOSS-26688	5420 Series rebooted unexpectedly VOSS 8.4.1.1 generating core file.
VOSS-26933	EDM Help does not open for the Configuration > Fabric > DVR > Globals tab.
VOSS-27040	Unexpected network congestion and traffic impact.
VOSS-27166	IP ECMP Max-Path 1 is configurable when ip ecmp enabled.
VOSS-27174	VSP 4900 Series 8.1.10.0 - ExtremeCloud IQ - Site Engine backup leads to I2C bus lockup reporting PSU, fan failures, and SFP reads incorrectly.
VOSS-27254	Failed to add new ARP records.
VOSS-27280	Reboot of DvR Leaf node with stp-multi-homing disables IS-IS and Leaf function.
VOSS-27419	Original DAC 10307 not working in 8.6.1.2.



Related Information

[MIB Changes](#) on page 108

MIB Changes

Deprecated MIBs

Table 33: Common

Object Name	Object OID	Deprecated in Release
rcIpBgpGeneralGroupRoutePolicyIn	1.3.6.1.4.1.2272.1.8.101.1.22	8.5
rcIpBgpGeneralGroupRoutePolicyOut	1.3.6.1.4.1.2272.1.8.101.1.23	8.5
rcIpConfOspfRfc1583Compatibility	1.3.6.1.4.1.2272.1.8.1.4.5	8.5

Modified MIBs

Table 34: Common

Object Name	Object OID	Modified in Release	Modification
rcChasType	1.3.6.1.4.1.2272.1.4.1	8.6.1	OTHER: Replace "VOSS" with "FabricEngine" in 5x20 models values
rc2kCardFrontType	1.3.6.1.4.1.2272.1.100.6.1.2	8.6.1	OTHER: Replace "VOSS" with "FabricEngine" in 5x20 models values

Table 34: Common (continued)

Object Name	Object OID	Modified in Release	Modification
rcLicenseLicenseType	1.3.6.1.4.1.2272.1.56.4	8.6.1	Added Enum:l10G4P(16), premierPlus10G4P(17), premierPlusMacsecPlus10G4P(18), macsecPlus10G4P(19), l10G8P(20), l10G4PPlus10G8P(21), premierPlus10G8P(22), premierPlusMacsecPlus10G8P(23), macsecPlus10G8P(24), premierPlus10G4PPlus10G8P(25), macsecPlus10G4PPlus10G8P(26), premierPlusMacsecPlus10G4PPlus10G8P(27)
rcIsidGlobalNameUsedByType	1.3.6.1.4.1.2272.1.87.6.1.4	8.8	ADD ENUM: radius(20)
rcIsidServiceOriginBitMap	1.3.6.1.4.1.2272.1.87.2.1.10	8.8	ADD ENUM: radiusL2Vsn(9)
rcIsidInterfaceOriginBitMap	1.3.6.1.4.1.2272.1.87.5.1.10	8.8	ADD ENUM: radiusL2Vsn(9)
rcIsisSpbmIpStaticIsidMcastVsnIsid	1.3.6.1.4.1.2272.1.63.30.1.2	8.8	Modified OID. Is part of the table index now.
rcIsisSpbmIpStaticIsidMcastGroup	1.3.6.1.4.1.2272.1.63.30.1.3	8.8	Modified OID
rcIsisSpbmIpStaticIsidMcastSource	1.3.6.1.4.1.2272.1.63.30.1.4	8.8	Modified OID
rcDvrGlobalRole	1.3.6.1.4.1.2272.1.219.1.2	8.8	CHANGE_RANGE: Changed the range from 1..2 to 1..3
rcVlanOrigin	1.3.6.1.4.1.2272.1.3.2.1.81	8.9	ADD_NEW_VALUE: ztf(3)
rcMACSecIfCAName	1.3.6.1.4.1.2272.1.88.2.1.1	8.9	CHANGE_RANGE: Changed range from 5..15 to 5..16
rc2kBootConfigEnableFactoryDefaults Mode	1.3.6.1.4.1.2272.1.100.5.1.60	8.9	ADD_NEW_VALUES: Add values for additional factorydefaults options

Table 35: 5320 Series

Object Name	Object OID	Modified in Release	Modification
rcLacpGlobalSmltSysId	1.3.6.1.4.1.2272.1.53.1.13	8.6.1	Not supported on 5320 platform
rcIsisPlsbSmltBmac	1.3.6.1.4.1.2272.1.63.4.1.10	8.6.1	Not supported on 5320 platform
rcIsisPlsbSmltPeerSysId	1.3.6.1.4.1.2272.1.63.4.1.11	8.6.1	Not supported on 5320 platform
rcIpConfRsmltEnable	1.3.6.1.4.1.2272.1.8.1.1.1.21	8.6.1	Not supported on 5320 platform

Table 35: 5320 Series (continued)

Object Name	Object OID	Modified in Release	Modification
rcIpConfRsmIltTable	1.3.6.1.4.1.2272.1.8.1.11	8.6.1	Not supported on 5320 platform
rcIpRsmIltEdgeSupportEnable	1.3.6.1.4.1.2272.1.8.26.1.2	8.6.1	Not supported on 5320 platform
rcMltMltType	1.3.6.1.4.1.2272.1.17.10.1.12	8.6.1	Only NORMAL MLT supported on 5320 platform
rc2kBootConfigAdvancedFeatureBwReservation	1.3.6.1.4.1.2272.1.100.5.1.51	8.9	OTHER: Update comment and description for low(3) value, 5320 and 5420 support only low(3) value

Table 36: 5420 Series

Object Name	Object OID	Modified in Release	Modification
SnpxChassisType		8.6	ADD ENUM: m532048T8XE, m532048P8XE, m532024T8XE, m532024P8XE, m532016P4XE, m532016P4XEDC OTHER: Replace "Virtual Services" with "Extreme Networks Fabric Engine" and "VOSS" with "FabricEngine" in comments only for Universal Hardware OTHER: Rebranding for Universal Hardware: Change enum values from m552048TVOSS, m552048WVOSS, m552012MW36WVOSS, m552024TVOSS, m552024WVOSS, m552024XVOSS, m552048SEVOSS to m552048T, m552048W, m552012MW36W, m552024T, m552024W, m552024X, m552048SE
rcSysLocatorLED	1.3.6.1.4.1.2272.1.1.125	8.6	OTHER: Add 5520, 5420 and 5320 in description
rcChasType	1.3.6.1.4.1.2272.1.4.1	8.6	ADD ENUM: a532048T8XEVOSS, a532048P8XEVOSS, a532024T8XEVOSS, a532024P8XEVOSS, a532016P4XEVOSS, a532016P4XEDCVOSS
rcIpConfGlobalTcpAdjustMssEnable	1.3.6.1.4.1.2272.1.8.1.6.29	8.6	OTHER: Add 5320 in description

Table 36: 5420 Series (continued)

Object Name	Object OID	Modified in Release	Modification
rcIpConfGlobalTcpAdjustMssStatus	1.3.6.1.4.1.2272.1.8.1.6.30	8.6	OTHER: Add 5320 in description
rcIpConfGlobalTcpAdjustMssType	1.3.6.1.4.1.2272.1.8.1.6.31	8.6	OTHER: Add 5320 in description
rcIpConfGlobalTcpAdjustMssValue	1.3.6.1.4.1.2272.1.8.1.6.32	8.6	OTHER: Add 5320 in description
rcIpfixAgingIntervalV2	1.3.6.1.4.1.2272.1.66.1.1.5	8.6	OTHER: Add 5320 in description
rc2kChassisPortLed3Status	1.3.6.1.4.1.2272.1.100.1.14	8.6	OTHER: Add 5320 in description
rc2kBootConfigEnableMacsec	1.3.6.1.4.1.2272.1.100.5.1.62	8.6	OTHER: Add 5320 in description
rc2kCardFrontType	1.3.6.1.4.1.2272.1.100.6.1.2	8.6	ADD ENUM: voss532048T8XE, voss532048P8XE, voss532024T8XE, voss532024P8XE, voss532016P4XE, voss532016P4XEDC
rc2kBootConfigAdvancedFeatureBwReservation	1.3.6.1.4.1.2272.1.100.5.1.51	8.9	OTHER: Update comment and description for low(3) value, 5320 and 5420 support only low(3) value

Table 37: 5520 Series

Object Name	Object OID	Modified in Release	Modification
rcPortAutoNegAd	1.3.6.1.4.1.2272.1.4.10.1.1.62	8.5	ADD_NEW_VALUE: advertise25000Full(13)
SnpxChassisType		8.6	ADD ENUM: m532048T8XE, m532048P8XE, m532024T8XE, m532024P8XE, m532016P4XE, m532016P4XEDC OTHER: Replace "Virtual Services" with "Extreme Networks Fabric Engine" and "VOSS" with "FabricEngine" in comments only for Universal Hardware
rcSysLocatorLED	1.3.6.1.4.1.2272.1.1.125	8.6	OTHER: Add 5520, 5420 and 5320 in description
rcChasType	1.3.6.1.4.1.2272.1.4.1	8.6	ADD ENUM: a532048T8XEVOSS, a532048P8XEVOSS, a532024T8XEVOSS, a532024P8XEVOSS, a532016P4XEVOSS, a532016P4XEDCVOSS
rcIpConfGlobalTcpAdjustMssEnable	1.3.6.1.4.1.2272.1.8.1.6.29	8.6	OTHER: Add 5320 in description
rcIpConfGlobalTcpAdjustMssStatus	1.3.6.1.4.1.2272.1.8.1.6.30	8.6	OTHER: Add 5320 in description

Table 37: 5520 Series (continued)

Object Name	Object OID	Modified in Release	Modification
rcIpConfGlobalTcpAdjustMssType	1.3.6.1.4.1.2272.1.8.1.6.31	8.6	OTHER: Add 5320 in description
rcIpConfGlobalTcpAdjustMssValue	1.3.6.1.4.1.2272.1.8.1.6.32	8.6	OTHER: Add 5320 in description
rcIpfixAgingIntervalV2	1.3.6.1.4.1.2272.1.66.1.1.5	8.6	OTHER: Add 5320 in description
rc2kChassisPortLed3Status	1.3.6.1.4.1.2272.1.100.1.14	8.6	OTHER: Add 5320 in description
rc2kBootConfigEnableMacsec	1.3.6.1.4.1.2272.1.100.5.1.62	8.6	OTHER: Add 5320 in description
rc2kCardFrontType	1.3.6.1.4.1.2272.1.100.6.1.2	8.6	ADD ENUM: voss532048T8XE, voss532048P8XE, voss532024T8XE, voss532024P8XE, voss532016P4XE, voss532016P4XEDC

Table 38: 5720 Series

Object Name	Object OID	Modified in Release	Modification
rcVossSystemCardLedId	1.3.6.1.4.1.2272.1.101.1.1.5.1.2	8.5	CHANGE_RANGE: Changed the range from 1..9 to 1..11
rcVossSystemTemperatureSensorIndex	1.3.6.1.4.1.2272.1.101.1.1.2.1.1	8.7	CHANGE_RANGE: Changed the range from 1..13 to 1..18
rcVxlanVtepSourceIp	1.3.6.1.4.1.2272.1.218.1	8.7	Not Supported on 5720
rcVxlanVtepVrf	1.3.6.1.4.1.2272.1.218.2	8.7	Not Supported on 5720
rcVxlanVtepTable	1.3.6.1.4.1.2272.1.218.3	8.7	Not Supported on 5720
rcVxlanVnidTable	1.3.6.1.4.1.2272.1.218.4	8.7	Not Supported on 5720
rcVossSystemVimAdminSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.3	8.7	ADD ENUM: mbps1000(4)
rc2kChassisPortLed3Status	1.3.6.1.4.1.2272.1.100.1.14	8.7	ADD_NEW_VALUES: Add values for speed and activity for 5720

New MIBs

Table 39: Common

Object Name	Object OID	New in VOSS Release
rcEapPortReauthOrigin	1.3.6.1.4.1.2272.1.57.2.1.29	8.6.1
rcEapPortReauthPeriodOrigin	1.3.6.1.4.1.2272.1.57.2.1.30	8.6.1
rcVossSystemFanInfoOperSpeedRpm	1.3.6.1.4.1.2272.1.101.1.1.4.1.6	8.6.1
rcVlanOrigin	1.3.6.1.4.1.2272.1.3.2.1.81	8.8
rcVlanMvnpnIsidValue	1.3.6.1.4.1.2272.1.3.2.1.82	8.8

Table 39: Common (continued)

Object Name	Object OID	New in VOSS Release
rcVlanMvvpnIsidStatus	1.3.6.1.4.1.2272.1.3.2.1.84	8.8
rcVrfOrigin	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.14	8.8
rcIpConfGlobalSpbMulticastPolicyEnable	1.3.6.1.4.1.2272.1.8.1.6.34	8.8
rcIpConfGlobalSpbMulticastPolicyRmap	1.3.6.1.4.1.2272.1.8.1.6.35	8.8
rcIpRoutePolicySetDatalsid	1.3.6.1.4.1.2272.1.8.100.13.1.49	8.8
rcIpRoutePolicySetRxOnly	1.3.6.1.4.1.2272.1.8.100.13.1.50	8.8
rcIpRoutePolicySetTxOnly	1.3.6.1.4.1.2272.1.8.100.13.1.51	8.8
rcIpConfGlobalSpbMulticastPolicyApply	1.3.6.1.4.1.2272.1.8.1.6.36	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastTable	1.3.6.1.4.1.2272.1.63.29.4	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastEntry	1.3.6.1.4.1.2272.1.63.29.4.1	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastType	1.3.6.1.4.1.2272.1.63.29.4.1.1	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastEnable	1.3.6.1.4.1.2272.1.63.29.4.1.2	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastIsidListName	1.3.6.1.4.1.2272.1.63.29.4.1.3	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastApply	1.3.6.1.4.1.2272.1.63.29.4.1.4	8.8
rcIsmMAL3RedistStaticIsidRoutedMcastRowStatus	1.3.6.1.4.1.2272.1.63.29.4.1.5	8.8
rcAutoSenseMultihostMacMax	1.3.6.1.4.1.2272.1.231.1.1.1.24	8.8
rcAutoSenseMultihostEapMacMax	1.3.6.1.4.1.2272.1.231.1.1.1.25	8.8
rcAutoSenseMultihostNonEapMacMax	1.3.6.1.4.1.2272.1.231.1.1.1.26	8.8
rcIsmCircuitPlsbl1MetricAuto	1.3.6.1.4.1.2272.1.63.5.1.10	8.9
rcPlugOptModSupportedSpeeds	1.3.6.1.4.1.2272.1.71.1.1.106	8.9
rcAutoSenseSisL1Metric	1.3.6.1.4.1.2272.1.231.1.1.1.27	8.9
rcAutoSenseSisL1MetricAuto	1.3.6.1.4.1.2272.1.231.1.1.1.28	8.9

Table 40: 5320 Series

Object Name	Object OID	New in Release
rc2kBootConfigEnableSpbmNodeScaling	1.3.6.1.4.1.2272.1.100.5.1.63	8.6
rcPrFilterAclStatsMatchDefaultPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.29	8.8
rcPrFilterAclStatsMatchDefaultPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.30	8.8

Table 40: 5320 Series (continued)

Object Name	Object OID	New in Release
rcPrFilterAclStatsMatchDefaultSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.31	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.32	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.33	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.34	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.35	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.36	8.8
rcVlanMvpngidOffset	1.3.6.1.4.1.2272.1.3.2.1.86	8.9
rcCloudIqLastDisconnectedTime	1.3.6.1.4.1.2272.1.230.11.24	8.9
rcCloudIqLastAttemptedAssociationTime	1.3.6.1.4.1.2272.1.230.11.25	8.9
rcCloudIqNextAssociationAttempt	1.3.6.1.4.1.2272.1.230.11.26	8.9
rcCloudIqAssociationFrequencyMode	1.3.6.1.4.1.2272.1.230.11.27	8.9

Table 41: 5420 Series

Object Name	Object OID	New in Release
rc2kBootConfigEnableSpbmNodeScaling	1.3.6.1.4.1.2272.1.100.5.1.63	8.6
rcPrFilterAclStatsMatchDefaultPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.29	8.8
rcPrFilterAclStatsMatchDefaultPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.30	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.31	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.32	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.33	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.34	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.35	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.11.2.3.2.1.36	8.8
rcVlanMvpngidOffset	1.3.6.1.4.1.2272.1.3.2.1.86	8.9

Table 41: 5420 Series (continued)

Object Name	Object OID	New in Release
rcCloudIqLastDisconnectedTime	1.3.6.1.4.1.2272.1.230.1.1.24	8.9
rcCloudIqLastAttemptedAssociationTime	1.3.6.1.4.1.2272.1.230.1.1.25	8.9
rcCloudIqNextAssociationAttempt	1.3.6.1.4.1.2272.1.230.1.1.26	8.9
rcCloudIqAssociationFrequencyMode	1.3.6.1.4.1.2272.1.230.1.1.27	8.9

Table 42: 5520 Series

Object Name	Object OID	New in Release
rcPrFilterAclStatsMatchDefaultPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.29	8.8
rcPrFilterAclStatsMatchDefaultPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.30	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.31	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.32	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.33	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.34	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.35	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.36	8.8
rcVlanMvnpnlsidOffset	1.3.6.1.4.1.2272.1.3.2.1.86	8.9
rcVossSystemAutoVimSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.8	8.9
rcCloudIqLastDisconnectedTime	1.3.6.1.4.1.2272.1.230.1.1.24	8.9
rcCloudIqLastAttemptedAssociationTime	1.3.6.1.4.1.2272.1.230.1.1.25	8.9
rcCloudIqNextAssociationAttempt	1.3.6.1.4.1.2272.1.230.1.1.26	8.9
rcCloudIqAssociationFrequencyMode	1.3.6.1.4.1.2272.1.230.1.1.27	8.9

Table 43: 5720 Series

Object Name	Object OID	New in Release
rcVossSystemVimGroupSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.7	8.7
rcPrFilterAclStatsMatchDefaultPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.29	8.8

Table 43: 5720 Series (continued)

Object Name	Object OID	New in Release
rcPrFilterAclStatsMatchDefaultPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.30	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.31	8.8
rcPrFilterAclStatsMatchDefaultSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.32	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.33	8.8
rcPrFilterAclStatsMatchGlobalPrimaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.34	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankPkts	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.35	8.8
rcPrFilterAclStatsMatchGlobalSecondaryBankOctets	1.3.6.1.4.1.2272.1.202.1.1.2.3.2.1.36	8.8
rcVlanMvpngIsidOffset	1.3.6.1.4.1.2272.1.3.2.1.86	8.9
rcVossSystemAutoVimSpeed	1.3.6.1.4.1.2272.1.101.1.1.8	8.9
rcCloudIqLastDisconnectedTime	1.3.6.1.4.1.2272.1.230.1.1.24	8.9
rcCloudIqLastAttemptedAssociationTime	1.3.6.1.4.1.2272.1.230.1.1.25	8.9
rcCloudIqNextAssociationAttempt	1.3.6.1.4.1.2272.1.230.1.1.26	8.9
rcCloudIqAssociationFrequencyMode	1.3.6.1.4.1.2272.1.230.1.1.27	8.9