



Fabric Engine Release Notes

For Fabric Engine Release 9.0.3

9037806-02 Rev AB
June 2024



Copyright © 2024 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: <https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks>

Open Source Declarations

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

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



Table of Contents

About this Document.....	6
Purpose.....	6
Conventions.....	6
Text Conventions.....	7
Documentation and Training.....	9
Open Source Declarations.....	9
Training.....	9
Help and Support.....	9
Subscribe to Product Announcements.....	10
Send Feedback.....	10
Document Revision Changes.....	12
New in this Release.....	13
New Software Features or Enhancements.....	13
General Enhancements.....	13
Multi-area SPB Enhancements.....	14
ExtremeCloud SD-WAN Enhancements.....	14
ZTP+ Enhancement.....	14
Other Changes.....	14
Scaling Updates.....	14
File Names for this Release.....	15
9.0.3 Feature Documentation.....	19
Multi-area SPB Concepts.....	20
Virtual NNI Links for Multi-Area Boundary Nodes.....	20
Multi-area SPB Considerations and Restrictions.....	21
Multi-area SPB CLI Tasks.....	23
Create a Virtual NNI Link Between Multi-Area Boundary Nodes.....	23
Multi-area SPB EDM Tasks.....	26
Create a Virtual NNI Link Between Multi-Area Boundary Nodes.....	27
Multi-area SPB Commands.....	30
ip address (loopback).....	30
logical-intf isis	31
show interfaces loopback.....	33
show isis logical-interface.....	34
ExtremeCloud SD-WAN Concepts.....	37
Fabric Extend (FE) States.....	37
Link Debounce.....	38
5320 Series VRF Support.....	39
ExtremeCloud SD-WAN CLI Tasks.....	40
Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration.....	40
Configure the IS-IS Area for a Specific Tunnel.....	41

Configure Auto-sense to Create All Learned Tunnels in the Remote Area.....	41
Display Auto-sense Configuration on the Switch.....	42
Extremecloud SD-WAN EDM Tasks.....	44
Specify the VRF for Auto-sense Extremecloud SD-WAN Configuration.....	44
Configure the IS-IS Area for a Specific Tunnel.....	47
Configure Auto-sense to Create All Learned Tunnels in the Remote Area.....	48
Extremecloud SD-WAN Commands.....	51
auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D} <home remote>.....	51
auto-sense sd-wan multi-area remote.....	52
auto-sense sd-wan vrf WORD<1-16>.....	52
link-debounce.....	53
show auto-sense.....	53
Other Documentation Changes.....	54
Default EDM Read Only Account.....	54
MLT Traffic Distribution Algorithm.....	54
Configure LLDP-MED Network Policies on Ports.....	55
Upgrade and Downgrade Considerations.....	57
Validated Upgrade Paths.....	57
Switches That Will Not Use Zero Touch Deployment.....	58
Switches That Will Use Zero Touch Deployment	58
Compatible Fabric IPsec Gateway Versions.....	60
Downgrade Considerations.....	60
Extremecloud IQ Agent.....	61
Downgrade Extremecloud IQ Managed Switches to 9.0.0.0.....	61
Post Upgrade Configuration for Zero Touch Fabric Configuration and Nickname Assignment.....	62
Network Requirements.....	63
Zero Touch Fabric Configuration Switch.....	63
Hardware and Software Compatibility.....	66
5320 Series Hardware.....	66
5420 Series Hardware.....	66
5520 Series Hardware.....	67
Operational Notes.....	68
Versatile Interface Module Operational Notes.....	69
Operational Notes for VIM Transceivers.....	69
5720 Series Hardware.....	69
Versatile Interface Module Operational Notes.....	70
7520 Series Hardware.....	70
7720 Series Hardware.....	71
Transceivers.....	71
Auto-Negotiation.....	71
Forward Error Correction (FEC).....	72
Power Supply Compatibility.....	72
Scaling.....	73
Layer 2.....	74
Maximum Number of Directed Broadcast Interfaces.....	79
Maximum Number of Microsoft NLB Cluster IP Interfaces.....	80
IP Unicast.....	80

IP Interface Maximums Clarification.....	90
IP Interface Maximums for 5320 Series.....	91
IP Interface Maximums for 5420 Series.....	91
IP Interface Maximums for 5520 Series.....	91
IP Interface Maximums for 5720 Series.....	92
IP Interface Maximums for 7520 Series.....	93
IP Interface Maximums for 7720 Series.....	93
Layer 3 Route Table Size.....	94
Route Scaling.....	94
IP Multicast.....	98
Distributed Virtual Routing (DvR).....	102
VXLAN Gateway.....	104
Filters, QoS, and Security.....	105
Filter Scaling.....	109
OAM and Diagnostics.....	117
Extreme Integrated Application Hosting Scaling.....	122
Fabric Scaling.....	123
Multi-area SPB Maximums.....	128
Number of I-SIDs Supported for the Number of Configured IS-IS Interfaces and Adjacencies.....	128
Interoperability Considerations for IS-IS External Metric.....	130
Recommendations.....	131
VRF Scaling.....	131
Important Notices.....	133
ExtremeCloud IQ Support.....	133
Compatibility with ExtremeCloud IQ Site Engine.....	133
Feature-Based Licensing	133
Memory Usage.....	134
Known Issues and Restrictions.....	135
Known Issues for this Release.....	135
Restrictions and Expected Behaviors.....	156
General Restrictions and Expected Behaviors.....	156
Filter Restrictions.....	163
Resolved Issues this Release.....	165
Related Information.....	166
MIB Changes.....	166
Deprecated MIBs.....	166
Modified MIBs.....	166
New MIBs.....	174
Obsolete MIBs.....	183



About this Document

[Purpose](#) on page 6

[Conventions](#) on page 6

[Documentation and Training](#) on page 9

[Help and Support](#) on page 9

[Send Feedback](#) on page 10

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.

Table 2: Text conventions

Convention	Description
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
Angle brackets (< >)	Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.

Table 3: Command syntax (continued)

Convention	Description
	If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code> , you can enter <code>cfm maintenance-domain maintenance-level 4</code> .
Bold text	Bold text indicates the GUI object name you must act upon. Examples: <ul style="list-style-type: none"> • Select OK. • On the Tools menu, choose Options.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. For example, if the command syntax is <code>ip address {A.B.C.D}</code> , you must enter the IP address in dotted, decimal notation.
Brackets ([])	Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is <code>show clock [detail]</code> , you can enter either <code>show clock</code> or <code>show clock detail</code> .
Ellipses (...)	An ellipsis (...) indicates that you repeat the last element of the command as needed. For example, if the command syntax is <code>ethernet/2/1 [<parameter> <value>] ...</code> , you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.
<i>Italic Text</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
Plain Courier Text	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: <ul style="list-style-type: none"> • <code>show ip route</code> • <code>Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]</code>

Table 3: Command syntax (continued)

Convention	Description
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation pane, expand Configuration > Edit .
Vertical Line ()	A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is <code>access-policy by-mac action { allow deny }</code> , you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code> , but not both.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

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

[Extreme Optics Compatibility](#)

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

Open Source Declarations

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

Training

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

Help and Support

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

[Extreme Portal](#)

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

[The Hub](#)

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

Call GTAC

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

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

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

Subscribe to Product Announcements

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

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

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

Send Feedback

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

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

To send feedback, email us at documentation@extremenetworks.com.

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



Document Revision Changes

The following table identifies changes between revisions of the same release document.

Table 4: 9.0.3 Release Notes revision changes

Revision	Change
AA	Initial revision for new release, see New in this Release on page 13
AB	Updated ECMP scaling in IP Unicast on page 80 for 7520 Series and 7720 Series



New in this Release

[New Software Features or Enhancements](#) on page 13

[Other Changes](#) on page 14

[File Names for this Release](#) on page 15

The following platforms support Fabric Engine 9.0.3:

- ExtremeSwitching 5320 Series
- ExtremeSwitching 5420 Series
- ExtremeSwitching 5520 Series
- ExtremeSwitching 5720 Series
- ExtremeSwitching 7520 Series
- ExtremeSwitching 7720 Series

For MIB-related changes, see [MIB Changes](#) on page 166.



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 Software Features or Enhancements

The following sections describe what is new in this release:

General Enhancements

This release introduces the following enhancements:

- Fail Open I-SID enhancement—You can now configure the Fail Open I-SID as the same I-SID value assigned by RADIUS VSA.
- LLDP-MED enhancement—You can now configure LLDP-MED network policies on ports using EDM. In previous releases, you could only view this information in EDM.
- RADIUS Dynamic Server—You can now configure up to eight clients.

Multi-area SPB Enhancements

This release adds the following Multi-Area enhancements for 7520 Series and 7720 Series:

- Increase the number of nodes that can function as boundary nodes from two to four.



Note

5520 Series and 5720 Series continue to support a maximum of two nodes that can function as boundary nodes.

- Ability to configure virtual NNI links for Multi-Area boundary nodes—Boundary nodes in the Multi-area SPB network require a robust Fabric path between them in both areas (home and remote). If a robust connection for one of the areas is not possible, you can create a virtual NNI link and establish a virtual Fabric adjacency over the area with the robust connection.

For more information, see [9.0.3 Feature Documentation](#) on page 19.

ExtremeCloud SD-WAN Enhancements

The software supports the following enhancements for ExtremeCloud SD-WAN:

- 5320 Series support— On models that support a single active VRF, you can now specify the VRF name that Auto-sense uses for the SD-WAN configuration.
- Auto-sense port Multi-area SPB support—On boundary nodes, you can configure in which IS-IS area Auto-sense creates an ExtremeCloud SD-WAN-learned interface.
- ExtremeCloud SD-WAN Bypass and MPLS support—Auto-sense automatically configures Link Debounce on the switch port that connects to SD-WAN Appliance. This configuration enables the switch that connects to the appliance LAN1 port to keep using its FE VXLAN tunnels over MPLS transport, even if SD-WAN Appliance is down, Layer 3 WAN Internet ports are lost, and the appliance is in Bypass mode.

For more information, see [9.0.3 Feature Documentation](#) on page 19.

ZTP+ Enhancement

In an earlier release, ZTP+ configuration supported assigning a CLIP in the GRT. Now the CLIP can be used for switch management.

Other Changes

Scaling Updates

[IP Unicast](#) on page 80 is updated to include DHCP client addresses.

File Names 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 file names 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 file names and sizes for this release.

Table 5: 5320 Series Software File names and Sizes

Description	File	Size
Logs reference	5320.9.0.3.0_edoc.tar	64,604,160 bytes
MD5 Checksum files	5320.9.0.3.0.md5	463 bytes
MIB - supported object names	5320.9.0.3.0_mib_sup.txt	1,550,250 bytes
MIB - objects in the OID compile order	5320.9.0.3.0_mib.txt	8,293,684 bytes
MIB - zip file of all MIBs	5320.9.0.3.0_mib.zip	1,234,445 bytes
Open source software - Master copyright file	5320.9.0.3.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	5320.9.0.3.0.sha512	1,378 bytes
Software image	5320.9.0.3.0.voss	116,727,264 bytes

Table 5: 5320 Series Software File names and Sizes (continued)

Description	File	Size
EDM Help files	FabricEnginev9.0.2_HELP_EDM_gzip.zip	5,235,518 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 6: 5420 Series Software File names and Sizes

Description	File	Size
Logs reference	5420.9.0.3.0_edoc.tar	64,604,160 bytes
MD5 Checksum files	5420.9.0.3.0.md5	463 bytes
MIB - supported object names	5420.9.0.3.0_mib_sup.txt	1,550,016 bytes
MIB - objects in the OID compile order	5420.9.0.3.0_mib.txt	8,293,684 bytes
MIB - zip file of all MIBs	5420.9.0.3.0_mib.zip	1,234,445 bytes
Open source software - Master copyright file	5420.9.0.3.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	5420.9.0.3.0.sha512	1,378 bytes
Software image	5420.9.0.3.0.voss	116,404,671 bytes
EDM Help files	FabricEnginev9.0.2_HELP_EDM_gzip.zip	5,235,518 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 7: 5520 Series Software File names and Sizes

Description	File	Size
Logs reference	5520.9.0.3.0_edoc.tar	64,604,160 bytes
MD5 Checksum files	5520.9.0.3.0.md5	463 bytes
MIB - supported object names	5520.9.0.3.0_mib_sup.txt	1,548,967 bytes
MIB - objects in the OID compile order	5520.9.0.3.0_mib.txt	8,293,684 bytes
MIB - zip file of all MIBs	5520.9.0.3.0_mib.zip	1,234,445 bytes
Open source software - Master copyright file	5520.9.0.3.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	5520.9.0.3.0.sha512	1,378 bytes
Software image	5520.9.0.3.0.voss	123,628,322 bytes

Table 7: 5520 Series Software File names and Sizes (continued)

Description	File	Size
EDM Help files	FabricEnginev9.0.2_HELP_EDM_gzip.zip	5,235,518 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 8: 5720 Series Software File names and Sizes

Description	File	Size
Logs reference	5720.9.0.3.0_edoc.tar	64,604,160 bytes
MD5 Checksum files	5720.9.0.3.0.md5	596 bytes
MIB - supported object names	5720.9.0.3.0_mib_sup.txt	1,555,725 bytes
MIB - objects in the OID compile order	5720.9.0.3.0_mib.txt	8,293,684 bytes
MIB - zip file of all MIBs	5720.9.0.3.0_mib.zip	1,234,445 bytes
Open source software - Master copyright file	5720.9.0.3.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	5720.9.0.3.0.sha512	1,703 bytes
Software image	5720.9.0.3.0.voss	334,446,870 bytes
EDM Help files	FabricEnginev9.0.2_HELP_EDM_gzip.zip	5,235,518 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

Table 9: 7520 Series Software File names and Sizes


Description	File	Size
Logs reference	7520.9.0.3.0_edoc.tar	64,604,160 bytes
MD5 Checksum files	7520.9.0.3.0.md5	596 bytes
MIB - supported object names	7520.9.0.3.0_mib_sup.txt	1,551,800 bytes
MIB - objects in the OID compile order	7520.9.0.3.0_mib.txt	8,293,684 bytes
MIB - zip file of all MIBs	7520.9.0.3.0_mib.zip	1,234,445 bytes
Open source software - Master copyright file	7520.9.0.3.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	7520.9.0.3.0.sha512	1,703 bytes
Software image	7520.9.0.3.0.voss	334,754,932 bytes
EDM Help files	FabricEnginev9.0.2_HELP_EDM_gzip.zip	5,235,518 bytes

Table 9: 7520 Series Software File names and Sizes (continued)

Description	File	Size
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

Table 10: 7720 Series Software File names and Sizes

Description	File	Size
Logs reference	7720.9.0.3.0_edoc.tar	64,604,160 bytes
MD5 Checksum files	7720.9.0.3.0.md5	596 bytes
MIB - supported object names	7720.9.0.3.0_mib_sup.txt	1,549,987 bytes
MIB - objects in the OID compile order	7720.9.0.3.0_mib.txt	8,293,684 bytes
MIB - zip file of all MIBs	7720.9.0.3.0_mib.zip	1,234,445 bytes
Open source software - Master copyright file	7720.9.0.3.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	7720.9.0.3.0.sha512	1,703 bytes
Software image	7720.9.0.3.0.voss	334,758,186 bytes
EDM Help files	FabricEnginev9.0.2_HELP_EDM_gzip.zip	5,235,518 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



9.0.3 Feature Documentation

- [Multi-area SPB Concepts](#) on page 20
- [Multi-area SPB CLI Tasks](#) on page 23
- [Multi-area SPB EDM Tasks](#) on page 26
- [Multi-area SPB Commands](#) on page 30
- [ExtremeCloud SD-WAN Concepts](#) on page 37
- [ExtremeCloud SD-WAN CLI Tasks](#) on page 40
- [ExtremeCloud SD-WAN EDM Tasks](#) on page 44
- [ExtremeCloud SD-WAN Commands](#) on page 51
- [Other Documentation Changes](#) on page 54

9.0.3 is a *Release Notes* only release. The topics in this section provide new or updated documentation for 9.0.3. For other feature information, see the 9.0.2 documentation suite:

- [Fabric Engine CLI Commands Reference](#)
- [Fabric Engine and VOSS Feature Support Matrix](#)
- [Fabric Engine User Guide](#)
- [Fabric Engine Alarms and Logs Reference](#)



Note

For Alarms and Logs information updated for 9.0.3, download the HTML files in the appropriate edoc.tar file. For more information, see [File Names for this Release](#) on page 15.

Multi-area SPB Concepts

The topics in this section provide conceptual-based documentation for new Multi-area SPB-related features.

Table 11: Multi-area SPB

Feature	Product	Release introduced
Multi-area SPB Boundary Node	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.4
Static data I-SID redistribution for Multi-area SPB Boundary Node	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 8.10
	5720 Series	Fabric Engine 8.10
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 8.8
Virtual NNI links for Multi-area SPB Boundary Nodes	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Not Supported
	5720 Series	Not Supported
	7520 Series	Fabric Engine 9.0.3
	7720 Series	Fabric Engine 9.0.3
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 9.0.3

Virtual NNI Links for Multi-Area Boundary Nodes

Boundary nodes in the Multi-area SPB network require a robust Fabric path between them in both areas (home and remote). If a robust connection for one of the areas is not possible, you can create a virtual NNI link and establish a virtual Fabric adjacency over the area with the robust connection. For instance, if an adjacency exists between two boundary nodes in the home area, you can use virtual NNI functionality to establish a Fabric adjacency in the remote area, and vice versa.

To create virtual NNI link functionality between boundary nodes, you must complete the following tasks:

- Configure the loopback IP address as the source IP address using the Multi-area virtual link flag.
- Configure a logical IS-IS interface using the Multi-area virtual link flag.
- Configure IS-IS on the logical interface.

IP Shortcuts automatically redistributes the IP address in the specific area without the need for an IS-IS redistribution policy. Boundary nodes receive this IP address in the corresponding area as an IS-IS redistributed route.

Virtual NNI Links on Boundary Nodes Considerations and Restrictions

The following list identifies considerations and restrictions that apply to virtual NNI links on boundary nodes:

- Function on boundary nodes only.
- Support only a CLIP (loopback) interface configured on the GRT.
- Support only one CLIP (loopback) interface.
- Do not support Bidirectional Forwarding Detection (BFD). You cannot enable BFD on a logical interface for a virtual NNI link.
- Cannot coexist with VXLAN Gateway. You cannot enable VXLAN Gateway on a loopback interface for a virtual NNI link and you cannot create a loopback interface for a virtual NNI link when VXLAN Gateway is enabled.

Multi-area SPB Considerations and Restrictions

The following list identifies the restrictions and considerations that apply to the Multi-area SPB feature:

- Two boundary nodes can be either in a vIST configuration (paired with each other) or in a non-vIST configuration. Three or more boundary nodes can only exist in a non-vIST configuration. Any other combination of boundary nodes is not supported.



Note

Only two boundary nodes can be in a vIST configuration.

- Up to four nodes can function as boundary nodes between any given pair of areas.



Note

5520 Series and 5720 Series support a maximum of two nodes that can function as boundary nodes.

- You must not connect the same Protocol Independent Multicast (PIM) domain to the SPB-PIM Gateway nodes that are in different Intermediate-System-to-Intermediate-System (IS-IS) areas, to avoid the inter-area redistribution of the same multicast information.
- You can enable the Dynamic Nickname server on the boundary nodes in the home area, but the boundary nodes cannot be clients in any of the two areas. The boundary nodes do not support the Dynamic Nickname server in the remote area.

- You must manually configure the backbone VLANs (B-VLAN) on the boundary nodes, so the system does not learn the dynamic values that it receives through the Link Layer Discovery Protocol (LLDP). However, the system sends the manually configured B-VLANs on the BN through LLDP, so that other neighbors can learn them (both in home and remote areas).
- Each time the port receives a Fabric Connect TLV, the port is configured as NNI in the home area. You must disable Auto-sense on the IS-IS remote area ports.
- If the system forms an adjacency between two boundary nodes that are part of the home and remote area, the hello packets in the home area use the home manual area and the hello packets in the remote area use the remote manual area.
- If the system forms a home and a remote adjacency on the same port then the Multi-area SPB feature uses different Backbone VLAN IDs (B-VIDs) for each adjacency, the home adjacency uses the primary B-VID and the remote adjacency uses the secondary B-VID.
- If the system forms an IS-IS adjacency in both the home and remote areas on a boundary node of the same port then the remote adjacency stays up only with another boundary node that also has IS-IS configured on both the home and remote areas of the same port.
- If a boundary node connects to a Backbone Edge Bridge (BEB) in the remote area and if you configure IS-IS in the home area on the same interface, then the remote adjacency goes down.
- On the boundary node, to install a route from a remote area in the routing table manager (RTM), the route must pass the accept policy and the Multi-area SPB redistribution policy that you configure on the specific Virtual Router Forwarding (VRF) instance.
- On the boundary node, to install an inter-VRF route from a remote area in the routing table manager (RTM), the inter-VRF route must pass the accept policy and the Multi-area SPB redistribution policy that you configure on both the source and destination VRF instances.
- Nickname and system ID for the physical node and virtual node must be different.
- When enabling Remote IS-IS Instance, make sure that the physical node nickname, virtual node nickname and system ID are different.
- You cannot establish an SSH connection to the boundary node from an IS-IS remote area.

Multi-area Deployment Guidelines

Use the following guidelines to design a Multi-area network with two or more Boundary nodes:

- In order to achieve optimal convergence performance, use the following order of preference for inter-connecting Boundary nodes:
 1. Direct links between all Boundary Nodes with adjacencies configured for both home and remote area.
 2. Reliable and redundant high bandwidth forwarding path between any two Boundary nodes in both areas.

3. Reliable and redundant high bandwidth forwarding path between any two Boundary nodes in home area and Virtual NNI Links in remote area.
 4. Reliable and redundant high bandwidth forwarding path between any two Boundary nodes in remote area and Virtual NNI Links in home area.
- Manually configure the Virtual NNI Links IS-IS metric to reflect the most desired traffic forwarding pattern. As a general rule, physical NNIs are preferred over Virtual NNI Links.
 - Manually configure the IS-IS metric of any relevant NNI in order to avoid using Virtual NNI Links to forward unicast and multicast traffic, when possible.

Multi-area SPB CLI Tasks

The topics in this section provide new or updated CLI task-based documentation related to Multi-area SPB support.

Create a Virtual NNI Link Between Multi-Area Boundary Nodes

Before You Begin

- Create basic SPBM and IS-IS infrastructure.
- Configure a CLIP interface.

About This Task

Perform the following procedure to create a virtual NNI link between Multi-area boundary nodes.

Procedure

1. Enter Loopback Interface Configuration mode


```
enable

configure terminal

interface Loopback <1-256>
```
2. Configure the loopback interface IP address to use as the source IP address for the Multi-area virtual NNI link:


```
ip address {<A.B.C.D/x> | <A.B.C.D> <A.B.C.D>} multi-area-virtual-link
<home | remote> [name WORD<0-64>]
```
3. Exit the Loopback Configuration mode:


```
exit
```
4. Create a IS-IS logical interface for the Multi-area virtual NNI link:


```
logical-intf isis <1-255> dest-ip <A.B.C.D> multi-area-virtual-link
[name WORD<1-64>]
```

5. Configure IS-IS on the logical interface:



Note

If you configure the loopback interface for the Multi-area virtual NNI link in the home area, you must configure IS-IS on the logical interface in the remote area, and vice versa.

a. Create an IS-IS interface:

```
isis [remote]
```

b. Enable the SPBM instance on the IS-IS interface:

```
isis [remote] spbm <1-100>
```

c. Enable the IS-IS interface:

```
isis [remote] enable
```

6. Verify the configuration using the following commands:

- show interfaces loopback
- show isis logical-interface

Example

Configure a Multi-area virtual link NNI for the home area. To do this, create the loopback source to be redistributed in the remote area and then enable IS-IS on a virtual-link logical interface in the home area.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 2.2.2.2/32 multi-area-virtual-link remote
exit
Switch:1(config)#logical-intf isis 2 dest-ip 4.4.4.4 multi-area-virtual-link
Switch:1(config-isis-2-4.4.4.4)#isis
Switch:1(config-isis-2-4.4.4.4)#isis spbm 1
Switch:1(config-isis-2-4.4.4.4)#isis enable
exit
```

Configure a Multi-area virtual NNI link for the remote area. To do this, create the loopback source to be redistributed in the home area and then enable IS-IS on a virtual-link logical interface in the remote area.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 1
Switch:1(config-if)#ip address 2.2.2.2/32 multi-area-virtual-link home
exit
Switch:1(config)#logical-intf isis 2 dest-ip 4.4.4.4 multi-area-virtual-link
Switch:1(config-isis-2-4.4.4.4)#isis remote
Switch:1(config-isis-2-4.4.4.4)#isis remote spbm 1
Switch:1(config-isis-2-4.4.4.4)#isis remote enable
exit
```

The following example displays loopback interface information:

```
Switch:1#show interfaces loopback

=====
                          Circuitless IP Interface - GlobalRouter
=====
INTF  IP_ADDRESS      NET_MASK      OSPF      PIM      AREA_ID      IF  IP      MULTI-AREA
```



```

=====
ID                STATUS  STATUS                INDX  NAME                VIRTUAL LINK
-----
1    2.2.2.2        255.255.255.255    disable  disable  0.0.0.0    1344  Virtual-link Remote
2    2.3.4.5        255.255.255.255    disable  disable  0.0.0.0    1345
15   192.0.2.2       255.255.255.0      disable  disable  0.0.0.0    1358  EXTR
=====
                                Loopback Ipv6 Interface
=====
IF   VRF      Descr      VLAN  PHYSICAL  ADMIN  OPER  TYPE  MTU  HOP  REACHABLE  RETRANSMIT  IPSEC
INDX NAME                                STATE  STATE                                LMT  TIME      TIME      STATE
-----
1234 GREEN    CLIPv6-11  00:00:00:00:00:0b  enable  up      ETHER 1500 64 30000  1000  disable
=====
                                Loopback IPv6 Address
=====
IPV6 ADDRESS/PREFIX LENGTH  LOOPBACK-ID  TYPE  ORIGIN  STATUS  VALID  PREF  NAME
                                LIFETIME  LIFETIME
-----
2001:DB8:2000::1/128        C-11         UNICAST  MANUAL  PREFERRED  INF  INF  EXTRSER200
=====

```

The following example displays logical interface information:

```

Switch:1>show isis logical-interface
=====
                                ISIS Logical Interfaces
=====
IFID  NAME      ENCAP  L2  INFO  VIDS  TUNNEL  L3  TUNNEL  NEXT  HOP  INFO  BFD  TUNNEL  ORIGIN  ISIS  SDWAN
ID    NAME      TYPE  PORT/MLT  (PRIMARY)  DEST-IP  PORT/MLT  VLAN  VRF      STATUS  SRC-IP  MTU  OPER  STATE
-----
1  SD-WAN-1  IP    --    --    --    192.0.2.3  Port1/44  4047  sd-wan  disabled  192.0.2.1  ZTF  1400  UP
2  SPBoIP_T1  IP    --    --    --    192.0.2.15  Port1/25  500  vrf23  disabled  192.0.2.16  CONFIG  1000  N/A
3  SPBoIP_T2  IP    --    --    --    192.0.2.224  MLT10  2  vrf24  disabled  192.0.2.22  CONFIG  1600  N/A
4  Virtual-link-4  IP    --    --    --    4.4.4.4  PortRX-NNI  4051  GlobalRouter  disabled  2.2.2.2  CONFIG  1600  N/A
5  Virtual-link-5  IP    --    --    --    5.5.5.5  Null  0  GlobalRouter  disabled  2.2.2.2  CONFIG  1600  N/A
=====
5 out of 5 Total Num of Logical ISIS interfaces
=====

```

Variable Definitions

The following table defines parameters for the **ip address** command.

Variable	Value
<A.B.C.D/X> <A.B.C.D> <A.B.C.D>	Specifies the IP address and subnet mask in the format A.B.C.D/X or A.B.C.D A.B.C.D.
multi-area-virtual-link	Specifies that the source IP address is used for a Multi-area virtual NNI link.
<home remote>	Specifies the transport area for the Multi-area virtual NNI link. Note: If you configure the loopback interface for the virtual NNI link in the home area, you must configure the IS-IS logical interface for the virtual NNI link in the remote area, and vice versa.
name WORD<0-64>	Specifies a name associated with the IP address. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as Virtual-link.

The following tables define parameters for the **logical-intf isis** command.

Variable	Value
<code><1-255></code>	Specifies the index number that uniquely identifies this logical interface.
<code>dest-ip <A.B.C.D></code>	Specifies the destination IP address of the Multi-area virtual NNI link. Note: This IP address is the loopback source IP address.
<code>multi-area-virtual-link</code>	Specifies that the logical interface is used for a Multi-area virtual NNI link.
<code>name WORD<1-64></code>	Specifies a name associated with the logical interface. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link-x</code> , where <code>x</code> is the logical interface ID.

The following tables define parameters for the **isis** command.

Variable	Value
<code>remote</code>	Specifies the remote area as the IS-IS interface. If the loopback interface for the virtual NNI link is configured in the remote area, omit this parameter.
<code>enable</code>	Enables the IS-IS interface for the virtual NNI link. The default is disabled.
<code>spbm <1-100></code>	Specifies the SPBM instance on the IS-IS interface.

Multi-area SPB EDM Tasks

The topics in this section provide new or updated EDM task-based documentation related to Multi-area SPB support.

Create a Virtual NNI Link Between Multi-Area Boundary Nodes

About This Task

Perform this procedure to create a virtual NNI link between Multi-area boundary nodes.

The assumption is that you are creating the CLIP and the logical interface for the first time. You can also use the table-based tab to apply the configuration to an existing interface.

Procedure

1. In the navigation pane, expand **Configuration > IP**.
2. Select **IP**.
3. Select the **Circuitless IP** tab.
4. Select **Insert**.
5. In the Interface field, add a CLIP interface number.
6. Type the IP address.
7. Type the network mask.
8. Select the type of loopback interface for the virtual NNI link:
 - **circuitlessIPMAVirtualLinkHome**
 - **circuitlessIPMAVirtualLinkRemote**



Note

If you configure the loopback interface for the virtual NNI link in the home area, you must configure IS-IS on the logical interface in the remote area, and vice versa.

9. (Optional) For **Name**, type a name for this interface.
10. Select **Insert**.
11. In the navigation pane, expand **Configuration > Fabric**.
12. Select **IS-IS**.
13. Select **Logical Interfaces** tab.
14. Select **Insert**.
15. For **Id**, type the index number that uniquely identifies this logical interface.
16. (Optional) For **Name**, type the name of this logical interface.
17. For **Type**, select **ip**.
18. For **DestIPAddr**, type the destination IP address for the logical interface.
19. Select **Multi-area Virtual Link** to configure a Multi-area virtual link on this interface.



Note

If you configure the loopback interface for the virtual NNI link in the home area, you must configure IS-IS on the logical interface in the remote area, and vice versa.

20. Select **Insert**.

What to Do Next

Configure IS-IS and SPBM on the logical interface.

Circuitless IP Field Descriptions

Use the data in the following table to use the **Circuitless IP** tab.

Name	Description
Interface	Specifies the number assigned to the interface.
Ip Address	Specifies the IP address of the CLIP.
Net Mask	Specifies the network mask.
Name	Specifies a name assigned to the IPv4 CLIP address. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link</code> .
IfType	Specifies the interface type. Note: Exception: circuitlessIPMAVirtualLinkHome and circuitlessIPMAVirtualLinkRemote apply to 7520 Series and 7720 Series only.

Logical Interfaces Field Descriptions

Use the data in the following table to use the **Logical Interfaces** tab and the **Insert Logical Interfaces** dialog. The available fields in the dialog differ depending on the type of core you select: **layer 2** or **ip**.

Name	Description
Id	Specifies the index number that uniquely identifies this logical interface. This field displays on the Insert Logical Interfaces dialog only.
IfIndex	Specifies the index number that uniquely identifies this logical interface. This field is read-only. This field displays on the Logical Interfaces tab only.

Name	Description
Name	Specifies a name associated with this logical interface, which can be up to 64 characters. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link-x</code> , where <code>x</code> is the logical interface ID.
Type	Specifies the type of logical interface to create: <ul style="list-style-type: none"> Specify layer 2 for a Layer 2 core network that the tunnel will traverse. Specify ip for a Layer 3 core network that the tunnel will traverse.
DestIPAddr	Specifies the destination IP address for the IP-type logical interface.
DestIfIndex	Specifies the physical port or MultiLink Trunking (MLT) that the Layer 2 logical interface is connected to.
Vids	Specifies the list of VLANs that are associated with this logical interface.
PrimaryVid	Specifies the primary tunnel VLAN ID associated with this Layer 2 Intermediate-System-to-Intermediate-System (IS-IS) logical interface.
CircIndex	Identifies the IS-IS circuit created under the logical interface. This field displays on the Logical Interfaces tab only.
NextHopVrf	Displays the next-hop VRF name to reach the logical tunnel destination IP. This field displays on the Logical Interfaces tab only. You can use this field to specify the VRF to reach the logical tunnel destination IP associated with a parallel tunnel.
ISIS Mtu	Specifies the Maximum Transmission Unit (MTU) size in bytes for IS-IS packets that use this logical interface. The default value is 1600.
BfdEnable	Enables or disables BFD on an IS-IS Logical Interface.
SrcIPAddr	Configures an additional source address to use as the parallel tunnel to create a backup adjacency. Note: To use an IPsec-encrypted tunnel as the parallel tunnel ensure that you configure the same source IP address on the logical IS-IS interface and in the Fabric IPsec Gateway virtual machine.

Name	Description
Origin	Specifies the origin of the IS-IS logical interface configuration, either through Zero Touch Fabric Configuration (ZTF) or manual configuration (config) through CLI or EDM.
Multi-Area Virtual Link Note: Exception: only applies to 7520 Series and 7720 Series.	Enables (true) or disables (false) a Multi-area boundary node virtual NNI link on an IS-IS Logical Interface. The default is disabled.

Multi-area SPB Commands

The topics in this section provide new or updated commands related to Multi-area SPB.

ip address (loopback)

Configure a circuitless IP interface (CLIP) when you want to provide a virtual interface that is not associated with a physical port. You can use a CLIP interface to provide uninterrupted connectivity to your switch.

Syntax

- `ip address <1-256> {A.B.C.D/X}`
- `ip address <1-256> {A.B.C.D/X} [vrf WORD<1-16>]`
- `ip address <1-256> {A.B.C.D/X} [name WORD<0-64>]`
- `ip address <1-256> {A.B.C.D} {A.B.C.D}`
- `ip address {A.B.C.D/X}`
- `ip address {A.B.C.D/X} [vrf WORD<1-16>]`
- `ip address {A.B.C.D/X} [name WORD<0-64>]`
- `ip address {A.B.C.D/X} multi-area-virtual-link <home | remote> [name WORD<0-64>]`
- `ip address {A.B.C.D} {A.B.C.D}`
- `ip address {A.B.C.D} {A.B.C.D} vrf WORD<1-16>`
- `ip address {A.B.C.D} {A.B.C.D} name WORD<0-64>`
- `ip address {A.B.C.D} {A.B.C.D} multi-area-virtual-link <home | remote> [name WORD<0-64>]`
- `no ip address <1-256> {A.B.C.D}`
- `no ip address <1-256> {A.B.C.D} vrf WORD<1-16>`
- `no ip address <1-256> {A.B.C.D} name WORD<0-64>`
- `no ip address {A.B.C.D}`
- `no ip address {A.B.C.D} vrf WORD<1-16>`
- `no ip address {A.B.C.D} name WORD<0-64>`

*Command Parameters***[vrf WORD<1-16>]**

Specifies an associated VRF by name.

{A.B.C.D/X}

Specifies the IP address and subnet mask.

{A.B.C.D}

Specifies the IP address.

<1-256>

Specifies the interface identification number for the circuitless IP (CLIP).

multi-area-virtual-link

Specifies that the source IP address is used for a Multi-area virtual NNI link.

<home | remote>

Specifies the transport area for the Multi-area virtual NNI link.

name WORD<0-64>

Specifies a name associated with the IP address.

**Note**

If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as `Virtual-link`.

Default

None

Command Mode

Loopback Interface Configuration

logical-intf isis

Create a logical IS-IS interface.

Syntax

- `logical-intf isis <1-255> dest-ip {A.B.C.D}`
- `logical-intf isis <1-255> dest-ip {A.B.C.D} name WORD<1-64>`
- `logical-intf isis <1-255> dest-ip {A.B.C.D} src-ip <A.B.C.D> [vrf WORD<1-16>]`
- `logical-intf isis <1-255> dest-ip {A.B.C.D} multi-area-virtual-link [name WORD<1-64>]`
- `logical-intf isis <1-255> vid {vlan-id[-vlan-id][, ...]} primary-vid <2-4059> mlt PT_MLT<1-512>`

- **logical-intf isis** <1-255> vid {vlan-id[-vlan-id][,...]} primary-vid <2-4059> port {slot/port[/sub-port]} name WORD<1-64>
- **no logical-intf isis** <1-255>

Command Parameters

<1-255>

Specifies the IS-IS logical interface ID.

dest-ip {A.B.C.D}

Specifies the destination IP address for the logical interface.

mIlt PT_MLT<1-512>

Specifies the MLT ID that the logical interface is connected to in a Layer 2 network.

multi-area-virtual-link

Specifies the logical IS-IS interface is for a virtual NNI link for Multi-area boundary nodes.

name WORD<1-64>

Specifies the administratively-assigned name of this logical interface.



Note

If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes for 7520 Series and 7720 Series, the system autogenerates the name and displays the name as `Virtual-link-x`, where `x` is the logical interface ID.

port {slot/port[/sub-port]}

Identifies a single slot and port. If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format `slot/port/sub-port`.

primary-vid <2-4059>

Specifies the primary tunnel VLAN ID associated with this Layer 2 IS-IS logical interface.

Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the `vrf-scaling` and `spbm-config-mode` boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.

src-ip <A.B.C.D> [vrf WORD<1-16>]

Configures an additional source address and optional VRF to use as the parallel tunnel for Fabric Extend.

To use an IPsec-encrypted tunnel as the parallel tunnel, ensure that you configure the same source IP address on the logical IS-IS interface and in the Fabric IPsec Gateway virtual machine.

vid {*vlan-id* [-*vlan-id*] [,...]}

Specifies the list of VLANs that are associated with this logical interface.

The VLAN ID is in one of the following formats: A single VLAN ID (*vlan-id*), a range of VLAN IDs [(*vlan-id*)-(*vlan-id*)] or a series of VLAN IDs (*vlan-id*, *vlan-id*, *vlan-id*).

Default

None.

Command Mode

Global Configuration

Usage Guidelines

The *multi-area-virtual-link* parameter only applies to 7520 Series and 7720 Series.

show interfaces loopback

Show loopback interface information.

Syntax

- **show interfaces loopback**
- **show interfaces loopback vrf** WORD <1-16> **name**
- **show interfaces loopback vrfids** WORD <0-512>

Command Parameters

name

Specifies the name associated with the IPv4 or IPv6 address.

vrfids WORD<0-512>

Specifies the ID of the VRF and is an integer in the range of 0 to 512.

vrf WORD<1-16>

Specifies the loopback information for the associated VRF name. WORD<1-16> specifies the VRF name in the range of 1 to 16 characters.

Default

None

Command Mode

Privileged EXEC

Example

```
Switch:1#show interfaces loopback
=====
Circuitless IP Interface - GlobalRouter
=====
INTF IP_ADDRESS NET_MASK OSPF PIM AREA_ID IF IP MULTI-AREA
ID STATUS STATUS AREA_ID INDX NAME VIRTUAL LINK
-----
1 2.2.2.2 255.255.255.255 disable disable 0.0.0.0 1344 Virtual-link Remote
2 2.3.4.5 255.255.255.255 disable disable 0.0.0.0 1345
15 192.0.2.2 255.255.255.0 disable disable 0.0.0.0 1358 EXTR
=====
Loopback Ipv6 Interface
=====
IF VRF Descr VLAN PHYSICAL ADMIN OPER TYPE MTU HOP REACHABLE RETRANSMIT IPSEC
INDX NAME ADDRESS STATE STATE LMT TIME TIME STATE
-----
1234 GREEN CLIPv6-11 00:00:00:00:00:0b enable up ETHER 1500 64 30000 1000 disable
=====
Loopback IPv6 Address
=====
IPv6 ADDRESS/PREFIX LENGTH LOOPBACK-ID TYPE ORIGIN STATUS VALID PREF NAME
LIFETIME LIFETIME
-----
2001:DB8:2000::1/128 C-11 UNICAST MANUAL PREFERRED INF INF EXTRSER200
=====
```

show isis logical-interface

Display IS-IS logical interfaces.

Syntax

- **show isis logical-interface [name]**

*Command Parameters***name**

Displays IS-IS logical interface name.

Default

None.

Command Mode

User EXEC

Command Output

The **show isis logical-interface** command displays the following information:

Output field	Description
IFIDX	Displays an index value for this logical interface.
NAME	Displays the name of this logical interface. Note: If you do not assign a name when you configure virtual NNI links for Multi-Area boundary nodes on 7520 Series and 7720 Series, the system autogenerates the name and displays the name as <code>Virtual-link-x</code> , where <code>x</code> is the loopback interface ID.
ENCAP TYPE	Displays whether the encapsulation type for the logical interface is Layer 2 (L2-P2P-VID) or Layer 3 (IP).
L2_INFO PORT/MLT	Displays the port or MLT that the logical interface is connected to in an Layer 2 network.
VIDS (PRIMARY)	Displays the list of VLANs that are associated with this Layer 2 logical interface.
TUNNEL DEST-IP	Displays the destination IP address for the logical interface.
L3_TUNNEL_NEXT_HOP_INFO PORT/MLT	Displays the outgoing interface (port or MLT) for VXLAN traffic.
L3_TUNNEL_NEXT_HOP_INFO VLAN	Displays the outgoing VLAN interface for VXLAN traffic.
L3_TUNNEL_NEXT_HOP_INFO VRF	Displays the name of the VRF that this Layer 3 logical interface is configured on.
BFD STATUS	Displays the status of BFD on this logical interface. The status can be enabled or disabled.
TUNNEL SRC-IP	Displays the source IP address for a Fabric Extend tunnel or a loopback interface.
ORIGIN	Displays the origin of the IS-IS logical interface configuration. For example, Zero Touch Fabric Configuration (ZTF) or manual configuration (config) through CLI or EDM.
ISIS MTU	Displays the Maximum Transmission Unit (MTU) size in bytes for IS-IS packets that use this logical interface. FE-IP deployments only.
SDWAN OPER STATE	Displays the one of the following states of the SD-WAN tunnel: <ul style="list-style-type: none"> • N/A (undefined) • UP • DOWN

Examples

Example of a Layer 2 Core (FE-VID):

```
Switch:1>show isis logical-interface
=====
ISIS Logical Interfaces
=====
IFIDX NAME      ENCAP  L2_INFO  VIDS    TUNNEL  L3_TUNNEL_NEXT_HOP_INFO  BFD  TUNNEL  ORIGIN  ISIS  SDWAN
TYPE           PORT/MLT (PRIMARY) DEST-IP  PORT/MLT  VLAN  VRF  STATUS  SRC-IP  MTU  OPER STATE
-----
1  SD-WAN-1 IP      --      --      192.0.2.3  Port1/44  4047  sd-wan  disabled  192.0.2.1  ZTF  1400  UP
2  --      L2-P2P-VID  Port2/1  101,201(101)  --      --      --      disabled  --      config  N/A
3  --      L2-P2P-VID  Port1/3  102,202(102)  --      --      --      disabled  --      config  N/A
-----
3 out of 3 Total Num of Logical ISIS interfaces
=====
```

Example of a Layer 3 Core (FE-IP):

```
Switch:1>show isis logical-interface
=====
ISIS Logical Interfaces
=====
IFIDX NAME      ENCAP  L2_INFO  VIDS    TUNNEL  L3_TUNNEL_NEXT_HOP_INFO  BFD  TUNNEL  ORIGIN  ISIS  SDWAN
TYPE           PORT/MLT (PRIMARY) DEST-IP  PORT/MLT  VLAN  VRF  STATUS  SRC-IP  MTU  OPER STATE
-----
1  SD-WAN-1 IP      --      --      192.0.2.3  Port1/44  4047  sd-wan  disabled  192.0.2.1  ZTF  1400  UP
2  SPBoIP_T1 IP      --      --      192.0.2.15  Port1/25  500  vrf23  disabled  192.0.2.16  CONFIG  1000  N/A
3  SPBoIP_T2 IP      --      --      192.0.2.224  MLT10  2  vrf24  disabled  192.0.2.22  CONFIG  1600  N/A
4  Virtual-link-4 IP      --      --      4.4.4.4  PortRX-NNI  4051  GlobalRouter  disabled  2.2.2.2  CONFIG  1600  N/A
5  Virtual-link-5 IP      --      --      5.5.5.5  Null  0  GlobalRouter  disabled  2.2.2.2  CONFIG  1600  N/A
-----
5 out of 5 Total Num of Logical ISIS interfaces
=====
```

The command **show isis logical-interface** truncates the IS-IS logical interface name to the first 16 characters. To view the entire name (up to a maximum of 64 characters), use the command **show isis logical-interface name**.

```
Switch:1>show isis logical-interface name
=====
ISIS Logical Interface name
=====
ID      NAME
-----
1      SD-WAN-1
2      SPBoIP_T1
3      SPBoIP_T2
4      Virtual-link-4
5      Virtual-link-5
-----
4 out of 4 Total Num of Logical ISIS interfaces
=====
```

ExtremeCloud SD-WAN Concepts

The topics in this section provide conceptual-based documentation for new ExtremeCloud SD-WAN-related features.

Table 12: SD-WAN product support

Feature	Product	Release introduced
Fabric Extend Integration with ExtremeCloud SD-WAN	5320 Series	Fabric Engine 8.10.1 Only 5320-48P-8XE and 5320-48T-8XE support more than one VRF with IP configuration. Beginning with Fabric Engine 9.0.3, you can specify the VRF that Auto-sense uses for SD-WAN on models that support a single active VRF.
	5420 Series	Fabric Engine 8.10.1
	5520 Series	Fabric Engine 8.10.1
	5720 Series	Fabric Engine 8.10.1
	7520 Series	Fabric Engine 8.10.1
	7720 Series	Fabric Engine 8.10.1
	VSP 4900 Series	VOSS 8.10.1
	VSP 7400 Series	VOSS 8.10.1
Auto-sense port Multi-area SPB support	5320 Series	Not Supported
	5420 Series	Not Supported
	5520 Series	Fabric Engine 9.0.3
	5720 Series	Fabric Engine 9.0.3
	7520 Series	Fabric Engine 9.0.3
	7720 Series	Fabric Engine 9.0.3
	VSP 4900 Series	Not Supported
	VSP 7400 Series	VOSS 9.0.3

The next section describes the Fabric Extend port state for Auto-sense.

Fabric Extend (FE) States

When Auto-sense is enabled, LLDP uses the FE TLV to create Fabric Extend tunnels between two Fabric switches that connect over the Internet through the SD-WAN Appliance. This functionality is supported on a single port of the switch.

The FE states are as follows:

- SD-WAN

- SD-WAN-PENDING

After the first Auto-sense port receives an FE-TLV, the port transitions to the SD-WAN state. All other Auto-sense ports transition to SD-WAN-PENDING state and remain unconfigured. When the first port transitions to the SD-WAN state, the switch verifies that VLAN 4047, VRF, and IS-IS logical interface configurations do not exist, and dynamically configures the following connectivity parameters:

- SD-WAN as the VLAN name associated with VLAN 4047 with origin ZTF
- sd-wan as the VRF name associated with the IP tunnel with origin DYNAMIC



Note

On switch models that support a single active VRF, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists but you can manually specify the VRF name that Auto-sense uses for the SD-WAN configuration on these models.

- SD-WAN-<ifidx> as the tunnel name
- SD-WAN Tunnel SrcIP as the name associated with the Fabric Extend underlay IP
- IPv4 address for VLAN 4047
- default route (0.0.0.0/0) with origin ZTF
- Fabric Extend tunnels with origin ZTF for IS-IS logical interfaces
- VLAN 4047 port membership
- Link Debounce timer of 8000 milliseconds on the switch port that connects to SD-WAN Appliance, if a timer configuration does not already exist

In the following cases, the port transitions to the SD-WAN-PENDING state:

- A secondary Auto-sense port receives an FE-TLV.
- The switch configuration includes the dynamic connectivity parameters, such as VLAN 4047, VRF, and IS-IS logical interfaces with the specified source IP address regardless of origin.

Link Debounce

Table 13: Link Debounce for WAN Links

Feature	Product	Release introduced
Link Debounce	5320 Series	Fabric Engine 8.6
	5420 Series	VOSS 8.5
	5520 Series	VOSS 8.5
	5720 Series	Fabric Engine 8.7
	7520 Series	Fabric Engine 8.10
	7720 Series	Fabric Engine 8.10
	VSP 4900 Series	VOSS 8.4.2
	VSP 7400 Series	VOSS 8.4.2

In a WAN environment, when a carrier-side link failure occurs, switchover on the carrier side can take a few hundred milliseconds. During that time, a lag in the sending and receiving of packets can occur. Use Link Debounce to hold the connection path until the switchover is complete. You can configure Link Debounce on each port.

Link Debounce protects the upper layers from unnecessary state changes by delaying the change of a port link state when the following situations occur:

- There are frequent flaps in a short interval at the physical layer in the case of Fiber WAN services.
- There is a delay in switching from the working path to the protected path in the case of Carrier Wave WAN services.

Link Debounce works only on Layer 1 protocol applications. Layer 2 / Layer 3 protocols make decisions based on how they receive packets. For example, STP makes the decision according to the lack of traffic and port up condition; OSPF and IS-IS can still fail adjacencies.

ExtremeCloud SD-WAN

Auto-sense automatically configures Link Debounce on the switch port that connects to SD-WAN Appliance. This configuration enables the switch that connects to the appliance LAN1 port to keep using its FE VXLAN tunnels over MPLS transport, even if SD-WAN Appliance is down, Layer 3 WAN Internet ports are lost, and the appliance is in Bypass mode.

If you do not configure a timer value and the port connects to SD-WAN Appliance, Auto-sense configures a value of 8000 milliseconds. Auto-sense does not overwrite a configured timer value.

5320 Series VRF Support

For the 5320 Series, only 5320-48P-8XE and 5320-48T-8XE support more than one VRF with IP configuration.

The 16- and 24-port 5320 Series models support a single active VRF with IP configuration; the VRF can be the Global Routing Table (GRT) or a non-default VRF:

- The GRT becomes the active VRF if you attach an IP interface to a VLAN. Otherwise, the first non-default VRF you create becomes the active VRF.
- If the GRT is the active VRF, you cannot create a second VRF.
- You cannot create more than one non-default VRF.
- You cannot configure IP interfaces in the GRT if the active VRF is a non-default VRF.
- To transition the active VRF from the GRT to a non-default VRF, remove all IP interfaces and then create the non-default VRF.
- To transition the active VRF from a non-default VRF to the GRT, first remove all router, ARP, and IP interfaces. Delete the VRF, and then add IP interfaces to the GRT.

Support for a single active VRF affects other features that rely on VRFs with IP configuration:

- ExtremeCloud SD-WAN—You must specify the VRF name that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.
- Fabric Extend—on platforms that support multiple VRFs with IP configuration, you cannot configure the tunnel source address (**ip-tunnel-source-address** command) if the address belongs to a VRF with an attached I-SID, which presents an issue for single VRF platforms. On platforms with only one active VRF, you can use an *overlay* parameter to bypass this restriction.
- Fabric Extend—you must configure a route-map policy to suppress IS-IS redistribution of the FE tunnel subnet:
 - Configure route-maps to not permit redistribution of the local route used as the tunnel source address (**ip-tunnel-source-address** command).
 - Configure an accept policy to deny IS-IS routes that overlap with the destination tunnel IP address.
- IP Shortcuts and Layer 3 VSN—You cannot configure the IP source address (**ip-source-address** command) as an IP address in the GRT if the active VRF is a non-default VRF.

ExtremeCloud SD-WAN CLI Tasks

The topics in this section provide new or updated CLI task-based documentation related to ExtremeCloud SD-WAN support.

Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration



Note

This procedure only applies to models that support a single active VRF. For more information, see [Fabric Engine and VOSS Feature Support Matrix](#).

Before You Begin

Ensure a user-defined VRF with a different name does not exist on the switch.

About This Task

Perform this procedure to specify the VRF that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.



Note

If the running switch configuration includes the dynamic SD-WAN VRF and you complete this procedure, the switch renames the dynamic VRF to your specified name, and converts the configuration to static.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Specify the VRF name that Auto-sense uses for ExtremeCloud SD-WAN configuration:

```
auto-sense sd-wan vrf WORD<1-16>
```

Configure the IS-IS Area for a Specific Tunnel

About This Task



Note

You can only use this command on a switch with boundary-node abilities.

Perform this procedure to specify the IS-IS area where Auto-sense creates a specific ExtremeCloud SD-WAN-learned tunnel. This configuration overrides a global area configuration for all learned tunnels.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```
2. Configure the IS-IS area:

```
auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D} <home | remote>
```

Configure Auto-sense to Create All Learned Tunnels in the Remote Area

About This Task



Note

You can only use this command on a switch with boundary-node abilities.

Perform this procedure to create all ExtremeCloud SD-WAN learned tunnels in the IS-IS remote area. You can also configure exceptions for specific tunnels. For more information, see [Configure the IS-IS Area for a Specific Tunnel](#) on page 41.

By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Procedure

1. Enter Global Configuration mode:

```
enable
```



```
configure terminal
```

2. Configure the IS-IS area:

```
auto-sense sd-wan multi-area remote
```

Display Auto-sense Configuration on the Switch

About This Task

Perform this procedure to display the Auto-sense configuration on the switch.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Display the Auto-sense configuration:

```
show auto-sense [access-differv] [data] [dhcp-detection] [eapol] [fa]
[isis] [onboarding] [qos] [sd-wan] [voice] [wait-interval]
```

3. Display the Auto-sense status and state on a port:

```
show interfaces gigabitEthernet auto-sense [{slot/port[/sub-port]}[-
slot/port[/sub-port]][,...]]
```

Examples

Display the Auto-sense status for Fabric Attach (FA):

```
Switch:1>show auto-sense fa
=====
                        AUTO-SENSE FA Config
=====
MSG-AUTH                MSG-AUTH-KEY
-----
enabled                  ****
-----
=====
                        AUTO-SENSE FA Client specific config
=====
TYPE                    EAPOL STATUS    I-SID  VLANID  C-VID  MGMT  I-SID  MGMT  C-VID
-----
camera                  Auto           100    100    untag  -     -     -
wap-type1               Auto           200    200    untag  -     -     -
open-virtual-switch    Auto           -      -      -      -     -     -
proxy-no-auth           Auth           300    300    untag  -     -     -
proxy                   Auth           400    n/a    400    400   400
-----
6 out of 6 Total Num of AUTO-SENSE entries displayed
=====
```

Display the Auto-sense configuration related to ExtremeCloud SD-WAN:

```
Switch:1>show auto-sense sd-wan
=====
                        AUTO-SENSE GLOBAL Config
=====
SDWAN
MULTI-AREA
-----
REMOTE
-----
=====
```

```

=====
                        AUTO-SENSE SD-WAN Logical Interfaces Multi-Area Config
=====
DEST-IP                MULTI-AREA
-----
10.10.10.10            REMOTE
20.20.20.20            HOME
30.30.30.30            REMOTE
-----
4 out of 4 Total Num of AUTO-SENSE entries displayed
=====

```

The following output displays an example from a switch that supports a single active VRF:

```

Switch:1>show auto-sense sd-wan
=====
                        AUTO-SENSE GLOBAL Config
=====
SDWAN
VRF
-----
Branch2
-----
1 out of 1 Total Num of AUTO-SENSE entries displayed
=====

```

Display the Auto-sense configuration related to voice:

```

Switch:1>show auto-sense voice
=====
                        AUTO-SENSE VOICE Config
=====
TYPE  LDDP-AUTH ENABLE I-SID      C-VID      DSCP      PRIORITY
-----
phone FALSE                2000       2000       46        6
-----
1 out of 1 Total Num of AUTO-SENSE entries displayed
=====

```

Display the global Auto-sense wait-interval information:

```

Switch:1>show auto-sense wait-interval
=====
                        AUTO-SENSE GLOBAL Config
=====
WAIT
INTERVAL
-----
50
-----
0 out of 0 Total Num of AUTO-SENSE entries displayed
=====

```

Display the Auto-sense status and state on a range of ports:

```

Switch:1>enable
Switch:1#show interfaces gigabitethernet auto-sense 1/1-1/5
=====
                        Port Auto-sense
=====

```

PORT NUM	AUTO-SENSE STATUS	AUTO-SENSE STATE	AUTO-SENSE PORT-DATA-ISID	AUTO-SENSE PORT-WAIT-INTERVAL
1/1	Enable	FA-PROXY-RING	--	20
1/2	Enable	SD-WAN	--	10
1/3	Enable	DOWN	--	10
1/4	Enable	DOWN	--	10
1/5	Disable	OFF	--	10

ExtremeCloud SD-WAN EDM Tasks

The topics in this section provide new or updated EDM task-based documentation related to ExtremeCloud SD-WAN support.

Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration



Note

This procedure only applies to models that support a single active VRF. For more information, see [Fabric Engine and VOSS Feature Support Matrix](#).

Before You Begin

Ensure a user-defined VRF with a different name does not exist on the switch.

About This Task

Perform this procedure to specify the VRF that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.



Note

If the running switch configuration includes the dynamic SD-WAN VRF and you complete this procedure, the switch renames the dynamic VRF to your specified name, and converts the configuration to static.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. In **Sd-Wan Vrf**, type the name of the VRF.
5. Select **Apply**.

Globals Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
AccessDiffservEnable	Enables or disables the differentiated service type as access for Auto-sense ports. The default is enabled.
Datalsid	Specifies the data I-SID used by the Auto-sense ports.
EapolVoiceLldpAuthEnable	Enables the EAPoL LLDP authentication for Auto-sense voice ports. The default is disabled.
FaMsgAuthEnable	Enables or disables the FA message authentication for Auto-sense ports. The default is enabled.
FaAuthenticationKey	Specifies the FA authentication key for Auto-sense ports.
IsisHelloAuthType	<p>Specifies the authentication type for IS-IS hello packets on Auto-sense ports:</p> <ul style="list-style-type: none"> • None • simple - simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5 - MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. • hmac-sha256 - with SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default authentication type is none.</p>
IsisHelloAuthKeyId	Specifies the IS-IS hello authentication number key id for the Auto-sense ports.
IsisHelloAuthKey	Specifies the IS-IS hello authentication number key for the Auto-sense ports. You must configure the IS-IS hello authentication key along with the IS-IS hello authentication type.
OnboardingIsid	Specifies the onboarding I-SID used by the Auto-sense ports.

Name	Description
Qos8021pOverrideEnable	Overrides the incoming 802.1p bits on ports that operate in Auto-sense mode. The default is enabled.
Voicelsid	Specifies the voice I-SID used by Auto-sense ports.
VoiceCvid	Specifies the customer VLAN ID associated with the voice I-SID used by Auto-sense ports. Voice C-Vid is configured for tagged voice traffic only. You must configure the Auto-sense voice customer VLAN ID along with the Auto-sense voice I-SID.
DhcpDetection	Enables or disables the DHCP detection in Auto-sense mode. The default is enabled.
FaCameraIsid	Specifies the FA camera I-SID used by Auto-sense ports.
FaProxyMgmtIsid	Specifies the FA proxy management I-SID used by Auto-sense ports.
FaProxyMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
FaProxyRingMgmtIsid	Specifies the FA proxy ring management I-SID used by Auto-sense ports.
FaProxyRingMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
FaProxyNoAuthIsid	Specifies the FA proxy no-auth I-SID used by Auto-sense ports.
FaVirtualSwitchIsid	Specifies the FA virtual-switch I-SID used by Auto-sense ports.
FaWapType1Isid	Specifies the FA WAP type-1 I-SID used by Auto-sense ports.
FaCameraEapolStatus	Specifies the FA EAPoL status for Camera I-SID used by Auto-sense ports.
FaEapolOVSStatus	Specifies the FA EAPoL status for OVS (Open-Virtual-Switch) I-SID used by Auto-sense ports.
FaEapolWap1Status	Specifies the FA EAPoL status for Wap-type-1 I-SID used by Auto-sense ports.
WaitInterval	Specifies the wait interval, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state. This configuration is a global configuration that applies to all Auto-sense ports. The default value is 35.
MultihostMacMax	Specifies the maximum number of EAPoL and non-EAPoL authentication MAC addresses allowed on this port. The default value is 2.
MultihostEapMacMax	Specifies the maximum number of EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port. The default value is 2.

Name	Description
MultihostNonEapMacMax	Specifies the maximum number of non-EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port. The default value is 2.
IsisL1Metric	Manually configure a value for the Level 1 metric. A higher number represents a higher cost and the least preferred route. The default value for L1 metric is 10 for any link, despite the port speed.
IsisL1MetricAuto	Enable the Level 1 metric as automatic. By enabling Level 1 metric as auto, the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed. This ensures that the fastest port speed determines the network route. The default is disabled.
FaProxyRingMgmtIsid	Specifies the FA proxy ring management I-SID used by Auto-sense ports.
FaProxyRingMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
Sd-Wan Area Note: This field only displays on a switch with boundary-node abilities.	Specifies the IS-IS area, home or remote, where Auto-sense creates all ExtremeCloud SD-WAN learned tunnels. By default, Auto-sense creates SD-WAN logical interfaces in the home area. You can also configure exceptions for specific tunnels. For more information, see Configure the IS-IS Area for a Specific Tunnel on page 47.
Sd-Wan Vrf Note: Exception: This field only applies to models that support a single active VRF. For more information, see Fabric Engine and VOSS Feature Support Matrix .	Specifies the VRF name that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.

Configure the IS-IS Area for a Specific Tunnel

About This Task



Note

You can only use this procedure on a switch with boundary-node abilities.

Perform this procedure to specify the IS-IS area where Auto-sense creates a specific ExtremeCloud SD-WAN-learned tunnel. This configuration overrides a global area configuration for all learned tunnels.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.

2. Select **AutoSense**.
3. Select the **Sd-Wan Multi-Area** tab.
4. Select **Insert**.
5. In **Ip**, type the destination IP address.
6. In **Area**, select **home** or **remote**.
7. Select **Insert**.

Sd-Wan Multi-Area *Field Descriptions*

Use the data in the following table to use the **Sd-Wan Multi-Area** tab.

Name	Description
Ip	Displays the destination IP address for an ExtremeCloud SD-WAN-learned tunnel.
Area	Specifies the IS-IS area where Auto-sense creates the tunnel. By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Configure Auto-sense to Create All Learned Tunnels in the Remote Area

About This Task



Note

You can only use this procedure on a switch with boundary-node abilities.

Perform this procedure to create all ExtremeCloud SD-WAN learned tunnels in the IS-IS remote area.

By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Procedure

1. In the navigation pane, expand **Configuration > Fabric**.
2. Select **AutoSense**.
3. Select the **Globals** tab.
4. In **Sd-Wan Area**, select **remote**.
5. Select **Apply**.

Globals *Field Descriptions*

Use the data in the following table to use the **Globals** tab.

Name	Description
AccessDiffservEnable	Enables or disables the differentiated service type as access for Auto-sense ports. The default is enabled.
Datalsid	Specifies the data I-SID used by the Auto-sense ports.
EapolVoiceLldpAuthEnable	Enables the EAPoL LLDP authentication for Auto-sense voice ports. The default is disabled.

Name	Description
FaMsgAuthEnable	Enables or disables the FA message authentication for Auto-sense ports. The default is enabled.
FaAuthenticationKey	Specifies the FA authentication key for Auto-sense ports.
IsisHelloAuthType	<p>Specifies the authentication type for IS-IS hello packets on Auto-sense ports:</p> <ul style="list-style-type: none"> • None • simple - simple password authentication uses a text password in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet. • hmac-md5 - MD5 authentication creates an encoded checksum in the transmitted packet. The receiving router uses an authentication key (password) to verify the MD5 checksum of the packet. • hmac-sha256 - with SHA-256 authentication, the switch adds an hmac-sha-256 digest to each Hello packet. The switch that receives the Hello packet computes the digest of the packet and compares it with the received digest. <p>Note: Secure Hashing Algorithm 256 bits (SHA-256) is a cipher and a cryptographic hash function of SHA2 authentication. You can use SHA-256 to authenticate IS-IS Hello messages. This authentication method uses the SHA-256 hash function and a secret key to establish a secure connection between switches that share the same key. This feature is in full compliance with RFC 5310.</p> <p>The default authentication type is none.</p>
IsisHelloAuthKeyId	Specifies the IS-IS hello authentication number key id for the Auto-sense ports.
IsisHelloAuthKey	Specifies the IS-IS hello authentication number key for the Auto-sense ports. You must configure the IS-IS hello authentication key along with the IS-IS hello authentication type.
OnboardingIsid	Specifies the onboarding I-SID used by the Auto-sense ports.
Qos8021pOverrideEnable	Overrides the incoming 802.1p bits on ports that operate in Auto-sense mode. The default is enabled.
VoiceIsid	Specifies the voice I-SID used by Auto-sense ports.
VoiceCvid	Specifies the customer VLAN ID associated with the voice I-SID used by Auto-sense ports. Voice C-Vid is configured for tagged voice traffic only. You must configure the Auto-sense voice customer VLAN ID along with the Auto-sense voice I-SID.
DhcpDetection	Enables or disables the DHCP detection in Auto-sense mode. The default is enabled.

Name	Description
FaCameraI Sid	Specifies the FA camera I-SID used by Auto-sense ports.
FaProxyMgmtI Sid	Specifies the FA proxy management I-SID used by Auto-sense ports.
FaProxyMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
FaProxyRingMgmtI Sid	Specifies the FA proxy ring management I-SID used by Auto-sense ports.
FaProxyRingMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
FaProxyNoAuthI Sid	Specifies the FA proxy no-auth I-SID used by Auto-sense ports.
FaVirtualSwitchI Sid	Specifies the FA virtual-switch I-SID used by Auto-sense ports.
FaWapType1I Sid	Specifies the FA WAP type-1 I-SID used by Auto-sense ports.
FaCameraEapOL Status	Specifies the FA EAPoL status for Camera I-SID used by Auto-sense ports.
FaEapOLOVSS Status	Specifies the FA EAPoL status for OVS (Open-Virtual-Switch) I-SID used by Auto-sense ports.
FaEapOLWap1 Status	Specifies the FA EAPoL status for Wap-type-1 I-SID used by Auto-sense ports.
WaitInterval	Specifies the wait interval, in seconds, for Auto-sense to wait for a Link Layer Discovery Protocol (LLDP) neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state. This configuration is a global configuration that applies to all Auto-sense ports. The default value is 35.
MultihostMacMax	Specifies the maximum number of EAPoL and non-EAPoL authentication MAC addresses allowed on this port. The default value is 2.
MultihostEapMacMax	Specifies the maximum number of EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port. The default value is 2.
MultihostNonEapMacMax	Specifies the maximum number of non-EAPoL authentication MAC addresses allowed on this port. Zero indicates that non-EAPoL authentication is disabled for this port. The default value is 2.
IsisL1Metric	Manually configure a value for the Level 1 metric. A higher number represents a higher cost and the least preferred route. The default value for L1 metric is 10 for any link, despite the port speed.

Name	Description
IsisLIMetricAuto	Enable the Level 1 metric as automatic. By enabling Level 1 metric as auto, the network route is determined by summing the lowest value metrics, which are inversely proportional to port speed. This ensures that the fastest port speed determines the network route. The default is disabled.
FaProxyRingMgmtIsid	Specifies the FA proxy ring management I-SID used by Auto-sense ports.
FaProxyRingMgmtCvid	Specifies the FA proxy management Client-VLAN ID (c-vid) used by Auto-sense ports.
Sd-Wan Area Note: This field only displays on a switch with boundary-node abilities.	Specifies the IS-IS area, home or remote, where Auto-sense creates all ExtremeCloud SD-WAN learned tunnels. By default, Auto-sense creates SD-WAN logical interfaces in the home area. You can also configure exceptions for specific tunnels. For more information, see Configure the IS-IS Area for a Specific Tunnel on page 47.
Sd-Wan Vrf Note: Exception: This field only applies to models that support a single active VRF. For more information, see Fabric Engine and VOSS Feature Support Matrix .	Specifies the VRF name that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.

ExtremeCloud SD-WAN Commands

The topics in this section provide new or updated commands related to ExtremeCloud SD-WAN.

`auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D} <home | remote>`

Specifies the area in which to create a specific ExtremeCloud SD-WAN learned tunnels. This configuration overrides a global area configuration for all learned tunnels.

Syntax

- `auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D} <home | remote>`
- `no auto-sense sd-wan multi-area logical-intf-dest-ip {A.B.C.D}`

Default

By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Command Mode

Global Configuration

Usage Guidelines

You can only use this command on a switch with boundary-node abilities.

auto-sense sd-wan multi-area remote

Creates all ExtremeCloud SD-WAN learned tunnels in the IS-IS remote area.

Syntax

- **auto-sense sd-wan multi-area remote**
- **no auto-sense sd-wan multi-area remote**

Default

By default, Auto-sense creates SD-WAN logical interfaces in the home area.

Command Mode

Global Configuration

Usage Guidelines

You can only use this command on a switch with boundary-node abilities.

auto-sense sd-wan vrf WORD<1-16>

Specifies the VRF name that Auto-sense uses for the SD-WAN configuration on models that support a single active VRF. On these models, Auto-sense cannot create the dynamic SD-WAN VRF if an IP configuration already exists.

Syntax

- **auto-sense sd-wan vrf WORD<1-16>**
- **no auto-sense sd-wan vrf WORD<1-16>**

Default

None.

Command Mode

Global Configuration

Usage Guidelines

This command does not apply to all hardware platforms. For more information about feature support, see [Fabric Engine and VOSS Feature Support Matrix](#).

link-debounce

Configure the Link Debounce timer for a port.

Syntax

- **default link-debounce**
- **link-debounce <0-300000>**
- **no link-debounce**

Command Parameters

<0-300000>

Specifies the Link Debounce time threshold in milliseconds.

Default

The default status is disabled for all ports when not initially configured. If you run the **default link-debounce** command, the default configuration is enabled with a value of 1,000 milliseconds. To return to the initial disabled state, you must run the **no link-debounce** command or set the Link Debounce timer to 0.

If you do not configure a timer value and the port connects to SD-WAN Appliance, Auto-sense configures a value of 8000 milliseconds.

Command Mode

GigabitEthernet Configuration.

Usage Guidelines

Auto-sense does not overwrite a configured timer value.

show auto-sense

Displays the Auto-sense configuration on the switch.

Syntax

- **show auto-sense [access-diffserv] [data] [dhcp-detection] [eapol] [fa] [isis] [onboarding] [qos] [sd-wan] [voice] [wait-interval]**

Command Parameters

access-diffserv

Displays the Auto-sense configuration related to Differentiated Services (DiffServ).

data

Displays the Auto-sense configuration related to the data I-SID.

dhcp-detection

Displays the Auto-sense configuration related to DHCP server auto-detection.

eapol

Displays the Auto-sense configuration related to Link Layer Discovery Protocol (LLDP) authentication for Extensible Authentication Protocol over LAN (EAPoL or EAP).

fa

Displays the Auto-sense configuration related to Fabric Attach (FA) message authentication and FA client-specific configuration.

isis

Displays the Auto-sense configuration related to Intermediate-System-to-Intermediate-System (IS-IS) authentication and information related to the L1 metric, such as a legend.

onboarding

Displays the Auto-sense configuration related to the onboarding I-SID.

qos

Displays the Auto-sense configuration related to overriding 802.1p bits.

sd-wan

Displays Auto-sense configuration related to ExtremeCloud SD-WAN.

voice

Displays the Auto-sense configuration related to voice for IP phones.

wait-interval

Displays the Auto-sense configuration related to the time to wait for an LLDP neighbor to be detected in the Auto-sense wait state before transitioning to the Auto-sense onboarding state.

Default

None.

Command Mode

User EXEC

Other Documentation Changes

The following sections provide smaller documentation updates.

Default EDM Read Only Account

The default user name for the EDM read-only account is user.

MLT Traffic Distribution Algorithm

The following table includes updated hash key information for Mac-In-Mac transit traffic. The hashing algorithm uses the following packet fields and the incoming

interface (source) port number to calculate the index to outgoing (destination) port number in an MLT:

Traffic type	Hashing algorithm
IPv4 traffic	Hash Key = [Destination IP Address (32 bits), Source IP Address (32 bits), Source TCP/UDP Port, Destination TCP/UDP port]
IPv4 traffic without TCP/UDP header	Hash Key = [Source IP Address (32 bits), Destination IP address (32 bits)]
IPv6 traffic	Hash Key = [Destination IPv6 Address (128 bits), Source IPv6 address (128 bits), Source TCP/UDP Port, Destination TCP/UDP port]
IPv6 traffic without TCP/UDP header	Hash Key = [Source IP Address (128 bits), Destination IP address (128 bits)]
Mac-In-Mac transit traffic	<ul style="list-style-type: none"> Inner Layer 2 non-IP packet: Hash Key = [Source Port (8bits), Customer VLAN(12bits), Customer Destination Mac Address (48 bits), Customer Source Mac Address (48 bits)] Inner IP packet: Hash Key = [Source Port (8bits), VLAN(12bits), Destination IP (32bits), Source IP (32bits)]
Layer 2 Non-IP traffic	Hash Key = [Destination MAC Address (48 bits), Source MAC Address(48 bits)]

Configure LLDP-MED Network Policies on Ports

About This Task

Perform this procedure to configure LLDP-MED network policies on specific ports.

Procedure

1. In the navigation pane, expand **Configuration > Serviceability > Diagnostics > 802_1ab**.
2. Select **Port MED**.
3. Select the **Local Policy** tab.
4. Select **Insert**.
5. In **PortNum**, select the ellipsis (...).
6. In **Port Editor: PortMembers** dialog box, select the desired ports.
7. Select **OK**.
8. In **PolicyAppType**, select the application type.
9. (Optional) In **PolicyVlanId**, type the VLAN ID for the port.
10. (Optional) In **PolicyPriority**, type the priority level.
11. (Optional) In **PolicyDscp**, type DSCP value.
12. (Optional) Select **Policy Tagged** to enable VLAN tagging on the port.
13. Select **Insert**.

Local Policy *Field Descriptions*

Use the data in the following table to use the **Local Policy** tab.

Name	Description
PortNum	Specifies the port.
PolicyAppType	Specifies the application type.
PolicyVlanId	Specifies the VLAN ID for the port, as defined in IEEE 802.1Q-2003. The value 0 is used if the device is using priority tagged frames, which means only the 802.1D priority level is significant, and the default VLAN ID of the ingress port is used instead.
PolicyPriority	Specifies the Layer 2 priority used for the specified application type, as defined in IEEE 802.1D-2004. The default is 0.
PolicyDscp	Specifies the value of the Differentiated Service Code Point (DSCP) associated with a specific port on the local LLDP-MED, as defined in IETF RFC 2474 and RFC 2475. The default is 0.
PolicyTagged	Specifies whether the application uses a tagged or untagged VLAN, as defined by IEEE 802.1Q-2003. <ul style="list-style-type: none"> • true — uses tagged VLAN • false — uses untagged VLAN or does not support a port-based VLAN



Upgrade and Downgrade Considerations

[Validated Upgrade Paths](#) on page 57

[Switches That Will Not Use Zero Touch Deployment](#) on page 58

[Switches That Will Use Zero Touch Deployment](#) on page 58

[Compatible Fabric IPsec Gateway Versions](#) on page 60

[Downgrade Considerations](#) on page 60

[Post Upgrade Configuration for Zero Touch Fabric Configuration and Nickname Assignment](#) on page 62

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:

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

- [Switches That Will Not Use Zero Touch Deployment](#) on page 58
- [Switches That Will Use Zero Touch Deployment](#) on page 58

Validated Upgrade Paths

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

**Note**

For any versions prior to 8.5.0.0, an intermediate upgrade is recommended because pre-8.5.0.0 versions are not validated.

Note that releases 8.6 and 8.7 are not validated upgrade paths. For non-validated upgrade paths, perform the upgrade with one or two switches initially before doing a widespread upgrade.

Table 14: Validated upgrade paths

Product	VOSS 8.5.x to Fabric Engine 9.0.x	Fabric Engine 8.8.x to Fabric Engine 9.0.x	Fabric Engine 8.9.x to Fabric Engine 9.0.x	Fabric Engine 8.10.x to Fabric Engine 9.0.x
5320 Series	N/A	Y	Y	Y
5420 Series	Y	Y	Y	Y
5520 Series	Y	Y	Y	Y
5720 Series	N/A	Y	Y	Y
7520 Series	N/A	N/A	N/A	Y
7720 Series	N/A	N/A	N/A	Y

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 57.
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 57.

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, 5720-48MXW, 7520 Series, and 7720 Series. 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 [File Names 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.



Caution

If you need to downgrade the image on ExtremeCloud IQ Managed Switches to release 9.0.0.0, from 9.0.2.0, or later, you must remove the file `.telegraf.csv` from the `/intflash` directory if it exists. Failure to do so can cause the switch to crash and revert to 9.0.2.0. For more information, see [Downgrade ExtremeCloud IQ Managed Switches to 9.0.0.0](#) on page 61.

ExtremeCloud IQ Agent

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](#).

Downgrade ExtremeCloud IQ Managed Switches to 9.0.0.0

Perform this procedure to downgrade switches that run GA version 9.0.2.0, or later, and are onboarded using ExtremeCloud IQ. This procedure does not apply to switches onboarded using ExtremeCloud IQ Site Engine.

Before You Begin

This procedure assumes the 9.0.0.0 GA image version is available on the switch. If not, you must upload it and extract the release distribution files to the `/intflash/release/` directory.

Procedure

1. Connect to the switch through the console, SSH, or Telnet.
2. Activate the 9.0.0.0 image:

```
enable

software activate 9.0.0.0 GA
```
3. Disable ExtremeCloud IQ Agent:

```
configure terminal

application

no iqagent enable
```
4. Delete the following file from the switch:

```
delete /intflash/.telegraf.csv -y
```
5. (Optional) Retain a copy of the current configuration, if needed:

```
copy config.cfg config.backup
```

6. Ensure the boot configuration points to the saved configuration from 9.0.0.0:

```
copy config.9.0.0.0 config.cfg
```

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

7. Reboot the switch to initiate the downgrade:

```
reset -y
```

8. Reconnect to the switch and commit the software:

```
enable
```

```
software commit
```

Post Upgrade Configuration for Zero Touch Fabric Configuration and 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.

For Fabric Engine 8.9, or earlier, 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 Zero Touch Fabric Configuration 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 devices running Fabric Engine 8.10 or later, the nickname automatically generates when you add new Zero Touch Fabric Configuration devices or implement Zero Touch Fabric Configuration on existing devices. You can configure a nickname server in your network with a dynamic nickname to replace the self-assigned nickname on your device.

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 57.

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:

- For devices running releases earlier than Fabric Engine 8.10, 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 159999999.
 - 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.

**Note**

This step only applies to devices running releases earlier than Fabric Engine 8.10.

5. The switch joins the Fabric.
6. For devices running releases earlier than Fabric Engine 8.10, the nickname server dynamically assigns an SPBM nickname. For devices running releases Fabric Engine 8.10, or later, the switch automatically assigns an SPBM nickname. The device searches the network for a nickname server and if one is found, the device replaces the automatic nickname with the dynamic nickname assigned by the server.
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 ExtremeCloud IQ Site Engine.



Hardware and Software Compatibility

- [5320 Series Hardware](#) on page 66
- [5420 Series Hardware](#) on page 66
- [5520 Series Hardware](#) on page 67
- [5720 Series Hardware](#) on page 69
- [7520 Series Hardware](#) on page 70
- [7720 Series Hardware](#) on page 71
- [Transceivers](#) on page 71
- [Power Supply Compatibility](#) on page 72

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 15: Switch models

Model	Initial Fabric Engine release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
5320-16P-4XE	8.6.1	Y	Y	Y	Y	Y
5320-16P-4XE-DC	8.6.1	Y	Y	Y	Y	Y
5320-24P-8XE	8.6	Y	Y	Y	Y	Y
5320-24T-8XE	8.6	Y	Y	Y	Y	Y
5320-48P-8XE	8.6	Y	Y	Y	Y	Y
5320-48T-8XE	8.6	Y	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 16: Switch models

Model	Initial release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
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 17: Switch models

Model	Initial release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
5520-12MW-36W	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-24T	AC: VOSS 8.2.5	Y - AC only	Y - AC only	Y	Y	Y
	ACDC: Fabric Engine 9.0					

Table 17: Switch models (continued)

Model	Initial release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
5520-24W	VOSS 8.2.5	Y	Y	Y	Y	Y
5520-24X	AC: VOSS 8.2.5	Y - AC only	Y - AC only	Y	Y	Y
	ACDC: Fabric Engine 9.0					
5520-48SE	AC: VOSS 8.2.5	Y - AC only	Y - AC only	Y	Y	Y
	ACDC: Fabric Engine 9.0					
5520-48T	AC: VOSS 8.2.5	Y - AC only	Y - AC only	Y	Y	Y
	ACDC: Fabric Engine 9.0					
5520-48W	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 18: Versatile Interface Modules (VIMs)

Model	Initial release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
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 supplies and fans with back-to-front airflow.
- The 5520-24T-ACDC, 5520-24X-ACDC, 5520-48SE-ACDC, and 5520-48T-ACDC models require a minimum of Fabric Engine 9.0 to support DC power supplies.

Versatile Interface Module Operational Notes

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

Table 19: 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
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 20: Switch models

Model	Initial Fabric Engine release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
5720-24MW	8.7	Y	Y	Y	Y	Y
5720-24MXW	8.7	Y	Y	Y	Y	Y

Table 20: Switch models (continued)

Model	Initial Fabric Engine release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
5720-48MW	8.7	Y	Y	Y	Y	Y
5720-48MXW	8.7	Y	Y	Y	Y	Y

**Note**

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

Table 21: Versatile Interface Modules (VIMs)

Model	Initial Fabric Engine release	Supported new Fabric Engine feature release				
		8.10	8.10.1	9.0	9.0.2	9.0.3
5720-VIM-2CE	8.7	Y	Y	Y	Y	Y
5720-VIM-6YE	8.7	Y	Y	Y	Y	Y

Versatile Interface Module Operational Notes

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

Table 22: 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

7520 Series Hardware

7520 Series is a universal hardware product that supports more than one Network Operating System (NOS) personality. For information about NOS personalities, see [ExtremeSwitching 7520 Series Hardware Installation Guide](#).

Table 23: Switch models

Model	Initial Fabric Engine release	Supported new Fabric Engine feature release			
		8.10.1	9.0	9.0.2	9.0.3
7520-48Y-8C	8.10	Y	Y	Y	Y
7520-48YE-8CE	9.0	N	Y	Y	Y
7520-48XT-6C	8.10	Y	Y	Y	Y

7720 Series Hardware

7720 Series is a universal hardware product that supports more than one Network Operating System (NOS) personality. For information about NOS personalities, see [ExtremeSwitching 7720 Series Hardware Installation Guide](#).

Table 24: Switch models

Model	Initial Fabric Engine release	Supported new Fabric Engine feature release			
		8.10.1	9.0	9.0.2	9.0.3
7720-32C	8.10	Y	Y	Y	Y

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](#)
- [ExtremeSwitching 7520 Series Hardware Installation Guide](#)
- [ExtremeSwitching 7720 Series Hardware Installation Guide](#)



Scaling

[Layer 2](#) on page 74

[IP Unicast](#) on page 80

[Layer 3 Route Table Size](#) on page 94

[IP Multicast](#) on page 98

[Distributed Virtual Routing \(DvR\)](#) on page 102

[VXLAN Gateway](#) on page 104

[Filters, QoS, and Security](#) on page 105

[OAM and Diagnostics](#) on page 117

[Extreme Integrated Application Hosting Scaling](#) on page 122

[Fabric Scaling](#) on page 123

[VRF Scaling](#) on page 131

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 25: Layer 2 Maximums

Attribute	Product	Maximum number supported
MAC table size (without SPBM)	5320 Series	32,000
	5420 Series	5420F Series models: 32,000 5420M Series models: 64,000
	5520 Series	80,000
	5720 Series	5720MXW models: 164,000 5720MW models: 100,000
	7520 Series	160,000
	7720 Series	160,000
MAC table size (with SPBM)	5320 Series	16,000
	5420 Series	5420F Series models: 16,000 5420M Series models: 32,000
	5520 Series	40,960
	5720 Series	5720MXW models: 82,000 5720MW models: 50,000
	7520 Series	80,000
	7720 Series	80,000
Endpoint Tracking MAC addresses per switch	5320 Series	n/a
	5420 Series	n/a
	5520 Series	8,000
	5720 Series	8,000
	7520 Series	8,000
	7720 Series	8,000

Table 25: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
Directed Broadcast interfaces	5320 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 79.
	5420 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 79.
	5520 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 79.
	5720 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 79.
	7520 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 79.
	7720 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 79.
Port-based VLANs Note: When you use Flex-UNI functionality, you can use the range from 1 to 4094 for port VLAN IDs.	5320 Series	4,059
	5420 Series	4,059
	5520 Series	4,059
	5720 Series	4,059
	7520 Series	4,059
	7720 Series	4,059
Private VLANs	5320 Series	See Table 26 on page 79
	5420 Series	See Table 26 on page 79
	5520 Series	See Table 26 on page 79
	5720 Series	See Table 26 on page 79
	7520 Series	See Table 26 on page 79
	7720 Series	See Table 26 on page 79

Table 25: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
Protocol-based VLANs (IPv6 only)	5320 Series	1
	5420 Series	1
	5520 Series	1
	5720 Series	1
	7520 Series	1
	7720 Series	1
RSTP instances	5320 Series	1
	5420 Series	1
	5520 Series	1
	5720 Series	1
	7520 Series	1
	7720 Series	1
MSTP instances	5320 Series	12
	5420 Series	12
	5520 Series	12
	5720 Series	12
	7520 Series	12
	7720 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
	7520 Series	56
	7720 Series	32 (up to 125 with channelization)

Table 25: 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
	7520 Series	8 active
	7720 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
	7520 Series	56
	7720 Series	32 (up to 125 with channelization)
Ports per MLT group	5320 Series	8 active
	5420 Series	8 active
	5520 Series	8
	5720 Series	8
	7520 Series	8
	7720 Series	8
Link State Tracking (LST) groups	5320 Series	48
	5420 Series	48
	5520 Series	48
	5720 Series	48
	7520 Series	48
	7720 Series	48

Table 25: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	8 upstream 128 downstream
	7720 Series	8 upstream 128 downstream
SLPP VLANs	5320 Series	128
	5420 Series	128
	5520 Series	128
	5720 Series	500
	7520 Series	500
	7720 Series	500
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.
	7520 Series	56
	7720 Series	32 (up to 125 with channelization)

Table 25: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
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 80.
	5720 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 80.
	7520 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 80.
	7720 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 80.

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 26: 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
7520 Series	100	50
7720 Series	100	50

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 27: 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 91.
	5420 Series	248 See IP Interface Maximums for 5420 Series on page 91.
	5520 Series	500 See IP Interface Maximums for 5520 Series on page 91.
	5720 Series	1,000 See IP Interface Maximums for 5720 Series on page 92.
	7520 Series	1,000 See IP Interface Maximums for 7520 Series on page 93
	7720 Series	1,000 See IP Interface Maximums for 7720 Series on page 93

Table 27: 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 91.
	5420 Series	124 See IP Interface Maximums for 5420 Series on page 91.
	5520 Series	252 See IP Interface Maximums for 5520 Series on page 91.
	5720 Series	500 See IP Interface Maximums for 5720 Series on page 92.
	7520 Series	500 See IP Interface Maximums for 7520 Series on page 93
	7720 Series	500 See IP Interface Maximums for 7720 Series on page 93
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 91.
	5520 Series	499 See IP Interface Maximums for 5520 Series on page 91.
	5720 Series	500 See IP Interface Maximums for 5720 Series on page 92.
	7520 Series	500 See IP Interface Maximums for 7520 Series on page 93
	7720 Series	500 See IP Interface Maximums for 7720 Series on page 93

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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 92.
	7520 Series	24 See IP Interface Maximums for 7520 Series on page 93
	7720 Series	24 See IP Interface Maximums for 7720 Series on page 93
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
	7520 Series	2,048/8
	7720 Series	2,048/8
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
	7520 Series	65
	7720 Series	65
OSPF v2/v3 neighbors (adjacencies)	5320 Series	50
	5420 Series	50
	5520 Series	100
	5720 Series	500
	7520 Series	500
	7720 Series	500

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	12 for each VRF 80 for the switch
	7720 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	5420F Series models: 15,000 5420M Series models: 24,000
	5520 Series	16,000
	5720 Series	5720MW Series models: 24,000 5720MXW Series models: 64,000
	7520 Series	40,000 with SPB
	7720 Series	40,000 with SPB
IPv4 CLIP interfaces	5320 Series	64
	5420 Series	64
	5520 Series	64
	5720 Series	64
	7520 Series	64
	7720 Series	64
IPv4 RIP interfaces	5320 Series	50
	5420 Series	50
	5520 Series	100
	5720 Series	200
	7520 Series	200
	7720 Series	200

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
IPv4 BGP peers	5320 Series	8
	5420 Series	8
	5520 Series	16
	5720 Series	256
	7520 Series	256
	7720 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
	7520 Series	16
	7720 Series	16
IPv4/IPv6 VRF instances For additional information, see VRF Scaling on page 131.	5320 Series	48-port models: 64 16- and 24-port models: 1 See IP Interface Maximums for 5320 Series on page 91.
	5420 Series	64 See IP Interface Maximums for 5420 Series on page 91.
	5520 Series	256 including mgmt VRF and GRT See IP Interface Maximums for 5520 Series on page 91.
	5720 Series	256 See IP Interface Maximums for 5720 Series on page 92.
	7520 Series	256 See IP Interface Maximums for 7520 Series on page 93
	7720 Series	256 See IP Interface Maximums for 7720 Series on page 93

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	2,000 for each VRF 10,000 for the switch
	7720 Series	2,000 for each VRF 10,000 for the switch
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
	7520 Series	1,000 for each VRF 5,000 for the switch
	7720 Series	1,000 for each VRF 5,000 for the switch

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	500 for each VRF 5,000 for the switch
	7720 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
	7520 Series	1,024
	7720 Series	1,024
DHCP client addresses provided by the DHCP server	5320 Series	1,000 clients
	5420 Series	10,000 clients
	5520 Series	10,000 clients
	5720 Series	100,000 clients
	7520 Series	100,000 clients
	7720 Series	100,000 clients
IPv4 DHCP Relay forwarding entries	5320 Series	248
	5420 Series	248
	5520 Series	512
	5720 Series	2,048
	7520 Series	2,048
	7720 Series	2,048

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	1,024
	7720 Series	1,024
IPv6 Neighbor table	5320 Series	8,000
	5420 Series	5420F Series models: 8,000 5420M Series models: 16,000
	5520 Series	16,000
	5720 Series	5720MW Series models: 24,000 5720MXW Series models: 32,000
	7520 Series	32,000
	7720 Series	32,000
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
	7520 Series	256
	7720 Series	256

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	128 per VRF 512 per system
	7720 Series	128 per VRF 512 per system
IPv6 CLIP interfaces	5320 Series	64
	5420 Series	64
	5520 Series	64
	5720 Series	64
	7520 Series	64
	7720 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
	7520 Series	1,000
	7720 Series	1,000
IPv6 6in4 configured tunnels	5320 Series	32
	5420 Series	32
	5520 Series	64
	5720 Series	64
	7520 Series	64
	7720 Series	64

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	512 per switch
	7720 Series	512 per switch
IPv6 BGP peers	5320 Series	8
	5420 Series	8
	5520 Series	16 Up to 8,000 IPv6 prefixes for BGPv6 peering
	5720 Series	256
	7520 Series	256
	7720 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
	7520 Series	16
	7720 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
	7520 Series	16
	7720 Series	16

Table 27: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	16
	7720 Series	16
BFD sessions per switch (IPv4) with 750ms timers for BGP and static routes only	5320 Series	48-port models: 16 16- and 24- port models: 1
	5420 Series	16
	5520 Series	16
	5720 Series	16
	7520 Series	50
	7720 Series	50
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
	7520 Series	16
	7720 Series	16

IP Interface Maximums Clarification

In the following sections, the formulas refer to "#IP Interfaces" count and not the count of IP addresses, which can be greater if you use IP multinetting with either IPv4 or IPv6. To clarify, if you use multinetting or IPv4 and IPv6 dual stack on a VLAN, the consumption of routable MAC resources is as follows:

- IPv4 address (primary) consumes one entry of routable MACs
- IPv4 address (primary) + any number of secondary addresses (multinetting) consumes one entry of routable MACs
- IPv6 interface (link-local) consumes one entry of routable MACs
- IPv6 interface (link-local) + any number of global addresses consume one entry of routable MACs
- IPv4 address (in any combination) + IPv6 interface (in any combination) consumes one entry of routable MACs

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

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

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

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

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 RSMLT 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 RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 3 = cannot exceed 248

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

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:
 - For interior node/non boundary node:
 - #NON DVR IP Interfaces + (# of VRRP interfaces) + (# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 3x(# of VRFs) + 1(if DVR node) + (#DVR VLANs if DVR controller) cannot exceed 1000
 - For boundary node:

#NON DVR IP Interfaces + 2x(# of VRRP interfaces) + 2x(# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 7x(# of VRFs) + 1(if DVR node) + 2x(#DVR VLANs if DVR controller) cannot exceed 1000

- If you enable the VRF scaling boot configuration flag:

- For interior node/non boundary node:

#NON DVR IP Interfaces + (# of VRRP interfaces) + (# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 3(if L3VSN is enabled) + 1(if DVR node) + (#DVR VLANs if DVR controller) cannot exceed 1000

- For boundary node:

#NON DVR IP Interfaces + 2x(# of VRRP interfaces) + 2x(# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 7(if L3VSN is enabled) + 1(if DVR node) + 2x(#DVR VLANs if DVR controller) cannot exceed 1000

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

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:

- For interior node/non boundary node:

#NON DVR IP Interfaces + (# of VRRP interfaces) + (# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 3x(# of VRFs) + 1(if DVR node) + (#DVR VLANs if DVR controller) cannot exceed 1000

- For boundary node:

#NON DVR IP Interfaces + 2x(# of VRRP interfaces) + 2x(# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 7x(# of VRFs) + 1(if DVR node) + 2x(#DVR VLANs if DVR controller) cannot exceed 1000

- If you enable the VRF scaling boot configuration flag:

- For interior node/non boundary node:

#NON DVR IP Interfaces + (# of VRRP interfaces) + (# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 3(if L3VSN is enabled) + 1(if DVR node) + (#DVR VLANs if DVR controller) cannot exceed 1000

- For boundary node:

#NON DVR IP Interfaces + 2x(# of VRRP interfaces) + 2x(# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 7(if L3VSN is enabled) + 1(if DVR node) + 2x(#DVR VLANs if DVR controller) cannot exceed 1000

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

IP Interface Maximums for 7520 Series

The maximum number of IP interfaces for 7520 Series is based on the following formulas:

- If you disable the VRF scaling boot configuration flag:
 - For interior node/non-boundary node:

$$\# \text{NON DVR IP Interfaces with unique mac offset} + (\# \text{ of VRRP interfaces}) + (\# \text{ of RSMLT interfaces}) + 2(\text{if IP Shortcuts is enabled}) + 3x(\# \text{ of VRFs}) + 1(\text{if DVR node}) + (\# \text{DVR VLANs if DVR controller})$$
 cannot exceed 1000
 - For boundary node:

$$\# \text{NON DVR IP Interfaces with unique mac offset} + 2x(\# \text{ of VRRP interfaces}) + 2x(\# \text{ of RSMLT interfaces}) + 2(\text{if IP Shortcuts is enabled}) + 7x(\# \text{ of VRFs}) + 1(\text{if DVR node}) + 2x(\# \text{DVR VLANs if DVR controller})$$
 cannot exceed 1000
- If you enable the VRF scaling boot configuration flag:
 - For interior node/non-boundary node:

$$\# \text{NON DVR IP Interfaces with unique mac offset} + (\# \text{ of VRRP interfaces}) + (\# \text{ of RSMLT interfaces}) + 2(\text{if IP Shortcuts is enabled}) + 3(\text{if L3VSN is enabled}) + 1(\text{if DVR node}) + (\# \text{DVR VLANs if DVR controller})$$
 cannot exceed 1000
 - For boundary node:

$$\# \text{NON DVR IP Interfaces with unique mac offset} + 2x(\# \text{ of VRRP interfaces}) + 2x(\# \text{ of RSMLT interfaces}) + 2(\text{if IP Shortcuts is enabled}) + 7(\text{if L3VSN is enabled}) + 1(\text{if DVR node}) + 2x(\# \text{DVR VLANs if DVR controller})$$
 cannot exceed 1000

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

IP Interface Maximums for 7720 Series

The maximum number of IP interfaces for 7720 Series is based on the following formulas:

- If you disable the VRF scaling boot configuration flag:
 - For interior node/non-boundary node:

$$\# \text{NON DVR IP Interfaces with unique mac offset} + (\# \text{ of VRRP interfaces}) + (\# \text{ of RSMLT interfaces}) + 2(\text{if IP Shortcuts is enabled}) + 3x(\# \text{ of VRFs}) + 1(\text{if DVR node}) + (\# \text{DVR VLANs if DVR controller})$$
 cannot exceed 1000
 - For boundary node:

$$\# \text{NON DVR IP Interfaces with unique mac offset} + 2x(\# \text{ of VRRP interfaces}) + 2x(\# \text{ of RSMLT interfaces}) + 2(\text{if IP Shortcuts is enabled}) + 7x(\# \text{ of VRFs}) + 1(\text{if DVR node}) + 2x(\# \text{DVR VLANs if DVR controller})$$
 cannot exceed 1000
- If you enable the VRF scaling boot configuration flag:
 - For interior node/non-boundary node:

#NON DVR IP Interfaces with unique mac offset + (# of VRRP interfaces) + (# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 3(if L3VSN is enabled) + 1(if DVR node) + (#DVR VLANs if DVR controller) cannot exceed 1000

- For boundary node:

#NON DVR IP Interfaces with unique mac offset + 2x(# of VRRP interfaces) + 2x(# of RSMLT interfaces) + 2(if IP Shortcuts is enabled) + 7(if L3VSN is enabled) + 1(if DVR node) + 2x(#DVR VLANs if DVR controller) cannot exceed 1000

For additional detail, see [IP Interface Maximums Clarification](#) on page 90.

Layer 3 Route Table Size

Table 28: Layer 3 Route Table Size Maximums

Attribute	Maximum number supported
IPv4 RIP routes	See Route Scaling on page 94.
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 29: 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 30: 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 31: 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 32: 5720 Series

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

Table 32: 5720 Series (continued)

URPF mode	IPv6 mode	5720 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
Yes	Yes	5720MW Series models: 4K 5720MXW Series models: 6K	5720MW Series models: 2K 5720MXW Series models: 3K	5720MW Series models: 1K 5720MXW 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).

Table 33: 7520 Series

URPF mode	IPv6 mode	7520 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
No	No	15,000	7,000	n/a
No	Yes	7,000	3,500	2,000
Yes	No	7,000	3,500	n/a
Yes	Yes	3,000	1,500	1,000

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 34: 7720 Series

URPF mode	IPv6 mode	7520 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
No	No	15,000	7,000	n/a
No	Yes	7,000	3,500	2,000
Yes	No	7,000	3,500	n/a

Table 34: 7720 Series (continued)

URPF mode	IPv6 mode	7520 Series		
		IPv4	IPv6 (prefix less than 64)	IPv6 (prefix greater than 64)
Yes	Yes	3,000	1,500	1,000

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 35: 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
	7520 Series	4,059
	7720 Series	4,059
PIM interfaces (IPv4/IPv6)	5320 Series	16 active
	5420 Series	16 active
	5520 Series	128 active
	5720 Series	128 active
	7520 Series	128 active
	7720 Series	128 active
PIM Neighbors (IPv4/IPv6) (GRT Only)	5320 Series	16
	5420 Series	16
	5520 Series	128
	5720 Series	128
	7520 Series	128
	7720 Series	128

Table 35: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
PIM-SSM static channels (IPv4/IPv6)	5320 Series	512
	5420 Series	512
	5520 Series	4,000
	5720 Series	4,000
	7520 Series	4,000
	7720 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
	7520 Series	6,000
	7720 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
	7520 Series	6,000
	7720 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
	7520 Series	3,000
	7720 Series	3,000
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
	7520 Series	4,000
	7720 Series	4,000

Table 35: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	2,000
	7720 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
	7520 Series	256
	7720 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
	7520 Series	6,000
	7720 Series	6,000
SPB-PIM Gateway controllers per SPB fabric (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	5
	5720 Series	n/a
	7520 Series	5
	7720 Series	5
SPB-PIM Gateway nodes per SPB fabric (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	64
	5720 Series	n/a
	7520 Series	64
	7720 Series	64

Table 35: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
SPB-PIM Gateway interfaces per BEB (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	64
	5720 Series	n/a
	7520 Series	64
	7720 Series	64
PIM neighbors per SPB-PIM Gateway node (IPv4)	5320 Series	n/a
	5420 Series	n/a
	5520 Series	64
	5720 Series	n/a
	7520 Series	64
	7720 Series	64

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 80.

Table 36: DvR Maximums

Attribute	Product	Maximum number supported
<p>Note:</p> <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. For VSP 4450 Series scaling information, see the VOSS Release Notes for VOSS Release 8.10. 		
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 250 on boundary node
	5720 Series	999 with vIST 1,000 without vIST 500 on boundary node
	7520 Series	999 with vIST as interior node 1,000 without vIST as interior node 500 on boundary node
	7720 Series	999 with vIST as interior node 1,000 without vIST as interior node 500 on boundary node
DvR domains per SPB fabric	5320 Series	16
	5420 Series	16
	5520 Series	16
	5720 Series	16
	7520 Series	16
	7720 Series	16

Table 36: 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
	7520 Series	8
	7720 Series	8
Leaf nodes per DvR domain	5320 Series	250
	5420 Series	250
	5520 Series	250
	5720 Series	250
	7520 Series	250
	7720 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 250 on boundary nodes
	5720 Series	999 with vIST 1,000 without vIST 500 on boundary nodes
	7520 Series	999 with vIST as interior node 1,000 without vIST as interior node 500 on boundary node
	7720 Series	999 with vIST as interior node 1,000 without vIST as interior node 500 on boundary node

Table 36: DvR Maximums (continued)

Attribute	Product	Maximum number supported
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	5420F Series models: 16,000 5420M Series models: 32,000
	5520 Series	48,000
	5720 Series	5720MW Series models: 64,000 5720MXW Series models: 96,000
	7520 Series	40,000
	7720 Series	40,000

VXLAN Gateway

Table 37: VXLAN Gateway Maximums

Attribute	Product	Maximum number supported
MAC addresses in base interworking mode	5320 Series	n/a
	5420 Series	n/a
	5520 Series	n/a
	5720 Series	n/a
	7520 Series	80,000
	7720 Series	80,000
MAC addresses in full interworking mode	5320 Series	n/a
	5420 Series	n/a
	5520 Series	n/a
	5720 Series	n/a
	7520 Series	50,000
	7720 Series	50,000
VNI IDs per node	5320 Series	n/a
	5420 Series	n/a
	5520 Series	n/a
	5720 Series	n/a
	7520 Series	2,000
	7720 Series	2,000

Table 37: VXLAN Gateway Maximums (continued)

Attribute	Product	Maximum number supported
VTEP destinations per node or VTEP	5320 Series	n/a
	5420 Series	n/a
	5520 Series	n/a
	5720 Series	n/a
	7520 Series	500
	7720 Series	500

The following table provides maximum numbers for OVSDB protocol support for VXLAN Gateway.

Table 38: OVSDB protocol support for VXLAN Gateway Maximums

Attribute	Product	Maximum number supported
Maximum controllers to which a single VTEP switch can connect	5320 Series	n/a
	5420 Series	n/a
	5520 Series	n/a
	5720 Series	n/a
	7520 Series	3
	7720 Series	3

Filters, QoS, and Security

For more information, see [Filter Scaling](#) on page 109.

Table 39: Filters, QoS, and Security Maximums

Attribute	Product	Maximum number supported
Total IPv4 Ingress rules/ ACEs (Port/VLAN/InVSN based, Security/QoS filters)	5320 Series	48-port models: 3,072 16- and 24-port models: 1,024
	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
	7520 Series	Primary Bank: 1,536 Secondary Bank: 1,536
	7720 Series	Primary Bank: 1,536 Secondary Bank: 1,536

Table 39: Filters, QoS, and Security Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	783 271 if you enable boot config flags ipv6-egress-filter
	7720 Series	783 271 if you enable boot config flags ipv6-egress-filter
Total IPv6 Ingress rules/ ACEs (Port/VLAN/InVSN based, Security filters)	5320 Series	1,024
	5420 Series	512
	5520 Series	512
	5720 Series	5720MW Series models: 1,536 5720MXW Series models: 2,048
	7520 Series	767
	7720 Series	767

Table 39: Filters, QoS, and Security Maximums (continued)

Attribute	Product	Maximum number supported
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	5720MW Series models: 1,536 5720MXW Series models: 3,072
	7520 Series	511
	7720 Series	511
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
	7520 Series	32
	7720 Series	32

Table 40: 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 boot config flags spbm-node-scaling: YES Platform VLAN: N/A

Table 40: NEAP Maximums (continued)

Product	Max # supported	Details
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	8,192	N/A
7520 Series	8,192	N/A
7720 Series	8,192	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:
 - (num ACLs + num ACEs + num of IPv4 Security Bank ACEs) <= 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:
 - (num ACLs + num ACEs) <=2982
- 4608 ingress ACEs

Ingress ACEs supported: (3072 Primary Bank - num ACLs) + (1536 Secondary Bank - num ACEs)

- 2982 egress ACEs

Egress ACEs supported: 2982 - 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:
 - ((num ACLs + num Primary Bank ACEs) <= 4096) && ((num ACLs + num Security Bank ACEs) <= 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:
 - (num ACLs + num ACEs + num of IPv4 Security Bank ACEs) <= 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:
 - $(\text{num ACLs} + \text{num ACEs}) \leq 6000$
- 6144 ingress ACEs

Ingress ACEs supported: $(4096 \text{ Primary Bank} - \text{num ACLs}) + (2048 \text{ Secondary Bank} - \text{num ACEs})$

- 6000 egress ACEs

Egress ACEs supported: $6000 - \text{num ACLs}$

7520 Series

The switch supports the following maximum limits for ACL scaling:

- 512 non-IPv6 ingress ACLs (inVSN, inPort, or inVlan):
 - 256 ACLs with 1 Primary ACE each + 256 ACLs with 1 Secondary ACE each OR
 - 383 ACLs with 1 Primary ACE each and/or 1 Secondary ACE each OR
 - a combination based on the following rule:
 - $\text{num ACLs} \leq 512 \ \&\& \ (\text{num ACLs} + \text{num Primary ACEs}) \leq 767 \ \&\& \ (\text{num ACLs} + \text{num Secondary ACEs}) \leq (767 - X)$ where $X = \text{num IPv6 ACLs} + \text{num IPv6 ACEs}$

For Primary bank, maximum implies a single port on inPort ACLs, a single I-SID for inVSN, and a single VLAN on inVlan ACLs.

For Secondary bank, inPort ACLs number of consumed rules is not multiplied by the number of ports attached to the ACL.

- 383 IPv6 ingress ACLs (inPort):
 - 383 IPv6 ACLs with 1 ACE each OR
 - A combination based on the following rule:
 - $\text{num IPv6 ACLs} \leq 383 \ \&\& \ (\text{num IPv6 ACLs} + \text{num ACEs}) \leq (767 - X)$ where $X = \text{num non-IPv6 ACLs} + \text{num non-IPv6 Secondary ACEs}$

This maximum implies a single port on inPort ACLs.

- 254 non-IPv6 egress ACLs (outPort):
 - 254 ACLs with 1 Security ACE each OR
 - A combination based on the following rule:
 - $\text{num ACLs} \leq 254 \ \&\& \ (\text{num ACLs} + \text{num Security ACEs}) \leq 508$

This maximum implies a single port on outPort ACLs.

- 256 IPv6 Egress ACLs (outPort):
 - 256 ACLs with 1 Security ACE each OR

- A combination based on the following rule:
 - $\text{num ACLs} \leq 256 \ \&\& \ (\text{num ACLs} + \text{num Security ACEs}) \leq 512$

This maximum implies a single port on outPort ACLs.

The switch supports the following maximum limits for ACE scaling:

- 1,532 non-IPv6 ingress ACEs

This theoretical maximum implies

- 2 non-IPv6 ingress ACL with 383+384 Primary ACEs and 383+384 Secondary ACEs
- no IPv6 ACLs configured
- a single port on inPort ACLs, and a single VLAN on inVLAN ACLs
- 767 IPv6 ingress ACEs

This theoretical maximum implies

- 1 IPv6 ingress ACL with 767 Security ACEs
- no non-IPv6 ACLs configured
- a port member count of 1 for inPort ACLs
- 783 non-IPv6 egress ACEs.

This theoretical maximum implies

- 1 egress ACL with 783 Security ACEs
- a port member count of 1 for outPort ACLs
- Non IPv6 egress ACEs supported: $783 - \text{num non-IPv6 egress ACLs}$
- 511 IPv6 egress ACEs

This theoretical maximum implies

- 1 egress ACL with 511 Security ACEs
- a port member count of 1 for outPort ACLs
- $511 - \text{num IPv6 egress ACLs}$

7720 Series

The switch supports the following maximum limits for ACL scaling:

- 512 non-IPv6 ingress ACLs (inVSN, inPort, or inVlan):
 - 256 ACLs with 1 Primary ACE each + 256 ACLs with 1 Secondary ACE each OR
 - 383 ACLs with 1 Primary ACE each and/or 1 Secondary ACE each OR
 - a combination based on the following rule:
 - $\text{num ACLs} \leq 512 \ \&\& \ (\text{num ACLs} + \text{num Primary ACEs}) \leq 767 \ \&\& \ (\text{num ACLs} + \text{num Secondary ACEs}) \leq (767 - X)$ where $X = \text{num IPv6 ACLs} + \text{num IPv6 ACEs}$

For Primary bank, maximum implies a single port on inPort ACLs, a single I-SID for inVSN, and a single VLAN on inVlan ACLs.

For Secondary bank, inPort ACLs number of consumed rules is not multiplied by the number of ports attached to the ACL.

- 383 IPv6 ingress ACLs (inPort):
 - 383 IPv6 ACLs with 1 ACE each OR
 - A combination based on the following rule:
 - $\text{num IPv6 ACLs} \leq 383 \ \&\& \ (\text{num IPv6 ACLs} + \text{num ACEs}) \leq (767 - X)$ where $X = \text{num non-IPv6 ACLs} + \text{num non-IPv6 Secondary ACEs}$

This maximum implies a single port on inPort ACLs.

- 254 non-IPv6 egress ACLs (outPort):
 - 254 ACLs with 1 Security ACE each OR
 - A combination based on the following rule:
 - $\text{num ACLs} \leq 254 \ \&\& \ (\text{num ACLs} + \text{num Security ACEs}) \leq 508$

This maximum implies a single port on outPort ACLs.

- 256 IPv6 Egress ACLs (outPort):
 - 256 ACLs with 1 Security ACE each OR
 - A combination based on the following rule:
 - $\text{num ACLs} \leq 256 \ \&\& \ (\text{num ACLs} + \text{num Security ACEs}) \leq 512$

This maximum implies a single port on outPort ACLs.

The switch supports the following maximum limits for ACE scaling:

- 1,532 non-IPv6 ingress ACEs

This theoretical maximum implies

- 2 non-IPv6 ingress ACL with 383+384 Primary ACEs and 383+384 Secondary ACEs
- no IPv6 ACLs configured
- a single port on inPort ACLs, and a single VLAN on inVLAN ACLs
- 767 IPv6 ingress ACEs

This theoretical maximum implies

- 1 IPv6 ingress ACL with 767 Security ACEs
- no non-IPv6 ACLs configured
- a port member count of 1 for inPort ACLs
- 783 non-IPv6 egress ACEs.

This theoretical maximum implies

- 1 egress ACL with 783 Security ACEs
- a port member count of 1 for outPort ACLs
- Non IPv6 egress ACEs supported: $783 - \text{num non-IPv6 egress ACLs}$
- 511 IPv6 egress ACEs

This theoretical maximum implies

- 1 egress ACL with 511 Security ACEs
- a port member count of 1 for outPort ACLs
- 511 - num IPv6 egress 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 41: 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
7520 Series	4	50	783	271
7720 Series	4	50	783	271

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 42: OAM and Diagnostics Maximums

Attribute	Product	Maximum number supported
EDM sessions	5320 Series	5
	5420 Series	5
	5520 Series	5
	5720 Series	5
	7520 Series	5
	7720 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)
	7520 Series	8 total (4 for IPv4 and 4 for IPv6)
	7720 Series	8 total (4 for IPv4 and 4 for IPv6)

Table 42: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
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)
	7520 Series	8 total (any combination of IPv4 and IPv6)
	7720 Series	8 total (any combination of IPv4 and IPv6)
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)
	7520 Series	16 total (8 for IPv4 and 8 for IPv6)
	7720 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)
	7520 Series	2 total (any combination of IPv4 and IPv6)
	7720 Series	2 total (any combination of IPv4 and IPv6)

Table 42: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	32 (up to 125 with channelization)
	7720 Series	32 (up to 125 with channelization)
Mirroring ports (destination)	5320 Series	4
	5420 Series	4
	5520 Series	4
	5720 Series	4
	7520 Series	4
	7720 Series	4

Table 42: 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.
	7520 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	7720 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.

Table 42: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
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.
	7520 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	7720 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
	7520 Series	1,000 Monitoring I-SIDs across SPB network
	7720 Series	1,000 Monitoring I-SIDs across SPB network
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
	7520 Series	3,100 samples per second
	7720 Series	3,100 samples per second

Table 42: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
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
	7520 Series	32,767
	7720 Series	32,767
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
	7520 Series	382 hosts
	7720 Series	382 hosts

Extreme Integrated Application Hosting Scaling

**Note**

The scaling attributes in this section apply to the following switches:

- 5720 Series models:
 - 5720-24MXW
 - 5720-48MXW
- 7520 Series
- 7720 Series

Table 43: Extreme Integrated Application Hosting (IAH) Maximums

Attribute	Product	Maximum number supported
Simultaneous Virtual Machines	5720-24MXW	2
	5720-48MXW	2
	7520 Series	6
	7720 Series	6
CPU cores available to VMs	5720-24MXW	2
	5720-48MXW	2
	7520 Series	6
	7720 Series	6

Table 43: Extreme Integrated Application Hosting (IAH) Maximums (continued)

Attribute	Product	Maximum number supported
Memory available to VMs	5720-24MXW	4 GB
	5720-48MXW	4 GB
	7520 Series	12 GB
	7720 Series	12 GB
Storage available to VMs	5720-24MXW	104 GB of 120 modular SSD
	5720-48MXW	104 GB of 120 modular SSD
	7520 Series	100 GB
	7720 Series	100 GB
Total SRIOV vports available to VMs	5720-24MXW	16
	5720-48MXW	16
	7520 Series	16
	7720 Series	16
Vports available to single VM	5720-24MXW	16
	5720-48MXW	16
	7520 Series	16
	7720 Series	16

Fabric Scaling

This section lists the fabric scaling information.

Table 44: 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 as boundary node	2
	5720 Series as boundary node	2
	7520 Series as boundary node	2
	7720 Series as boundary node	2

Table 44: Fabric maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
Number of B-VIDs	5320 Series	2
	5420 Series	2
	5520 Series	2
	5720 Series	2
	7520 Series	2
	7720 Series	2
Maximum number of Physical and Logical (Fabric Extend) NNI interfaces/adjacencies (Home and Remote area total when operating as Boundary Node)	5320 Series (cannot operate as boundary node)	64
	5420 Series (cannot operate as boundary node)	50
	5520 Series	128
	5720 Series	128
	7520 Series	255
	7720 Series	255
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
	7520 Series	See Number of I-SIDs supported
	7720 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
	7520 Series	4,000
	7720 Series	4,000

Table 44: Fabric maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
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
	7520 Series	56 (up to 125 with channelization)
	7720 Series	32 (up to 125 with channelization)
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
	7520 Series	100
	7720 Series	100
Maximum number of routed PVLANs/E-Trees	5320 Series	10
	5420 Series	10
	5520 Series	10
	5720 Series	100
	7520 Series	50
	7720 Series	50
Maximum number of Layer 3 VSNs per switch See VRF Scaling on page 131.	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
	7520 Series	256
	7720 Series	256

Table 44: Fabric maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
Maximum number of FA I-SID/ VLAN assignments per port	5320 Series	94
	5420 Series	94
	5520 Series	94
	5720 Series	94
	7520 Series	94
	7720 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
	7520 Series	50,000
	7720 Series	50,000
ISW switches in a Fabric Attach Ring		128

Table 45: Multidimensional Fabric node scale

Device	Node scaling ¹	SPBM nodes ²	Total unicast BMACs ³	Switched UNI endpoints ⁴	Multicast Data I-SIDs ⁵	
					Ingress BEB	Egress BEB
5320 Series	Enabled	500	500	400	16- and 24-port models: 250 48-port models: 500	1,200
	Disabled	350	350	700/800	16- and 24-port models: 250 48-port models: 500	800
5420 Series	Enabled	500 without vIST 340 with vIST	500 without vIST 340 with vIST	400	500	1,200
	Disabled	350 without vIST 340 with vIST	350 without vIST 340 with vIST	700/800	500	800
5520 Series		500/800	800	2,700	2,700	4,000
5720 Series		500/1,000	2,000	4,850	4,000	6,000
7520 Series		500/1,000	2,000	12,000	4,000	6,000

Table 45: Multidimensional Fabric node scale (continued)

Device	Node scaling ¹	SPBM nodes ²	Total unicast BMACs ³	Switched UNI endpoints ⁴	Multicast Data I-SIDs ⁵	
					Ingress BEB	Egress BEB
7720 Series		500/1,000	2,000	12,000	4,000	6,000

1. Node scaling—refers to the enabled state of the **boot configuration flags spbm-node-scaling** command, if applicable. For 5320 Series and 5420 Series using Fabric Extend or vIST, it is a best practice to enable this command.
2. SPBM nodes—refers to the number of supported SPBM enabled nodes, both BEB and BCB. When different, the number is formatted as per area/total per device. For 5420 Series, this number is impacted by vIST.
3. Total unicast BMACs—refers to the total number, both virtual and physical, this node can share services with. This number includes Layer 2 VSNS, Layer 3 VSNS, E-TREE, Multicast, and Transparent Port UNI. For 5420 Series, this number is impacted by vIST.
4. Switched UNI endpoints—refers to the maximum local tagged and untagged endpoints, either manual, RADIUS, or FA-assigned. When different, the number is formatted based on the configuration of the **boot config flags macsec** command: enabled/disabled.
5. Multicast Data I-SIDs—refers to the maximum Layer 2 or Layer 3, dynamic and static originated data I-SIDs. The overall limits are across all locally configured Layer 2 VSNS

The following table provides numbers for 5320 Series and 5420 Series only, to reflect the impact of the **boot configuration flags spbm-node-scaling** command.

Table 46: Maximum remote multicast sender nodes and local I-SIDs

Device	Node scaling ¹	Total remote multicast sender nodes ²	Total local I-SIDs ³
5320 Series	Enabled	200	16- and 24-port models: 274 48-port models: 500
	Disabled	150	16- and 24-port models: 274 48-port models: 564
5420 Series	Enabled	200	500
	Disabled	150	564

1. Node scaling—refers to the enabled state of the **boot configuration flags spbm-node-scaling** command. For 5320 Series and 5420 Series using Fabric Extend or vIST, it is a best practice to enable this command.
2. Total remote multicast sender nodes—refers to the total number of nodes that send IP multicast streams that the local BEB receives. This space is shared with unicast BMACs in the preceding table. Documented limits are individual in isolation; introducing vIST clusters or nodes that advertise IP multicast streams decreases the total number of physical nodes in an area.
3. Total local I-SIDs—refers to the total for Layer 2, Layer 3, and Multicast.

Multi-area SPB Maximums

Table 47: Multi-area SPB Maximums

Scaling	5520 Series	5720 Series	7520 Series	7720 Series
SPBM enabled nodes per area	500	500	500	500
SPBM total nodes home + remote	650	650	1,000	1,000
I-SIDs supported on boundary nodes (no local UNI present on device)	2,000	2,000	9,600	9,600
Maximum number of inter-area redistributed Layer 2 VSNs (no local UNI present on Boundary Node)	2,000	2,000	9,600	9,600
Maximum number of IP multicast S,Gs when operating as a boundary node (inter-area)	1,600	1,600	4,800	4,800
DvR host routes redistributed across area boundary	n/a	6,000	13,900	13,900
SPBM multicast-FIB entries	10,000	10,000	35,000	35,000

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

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
	7520 Series	4,000	4,000
	7720 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
	7520 Series	4,000	4,000
	7720 Series	4,000	4,000
10	5320 Series	n/a	500
	5420 Series	564	564
	5520 Series	2,900	4,000
	5720 Series	2,900	4,000
	7520 Series	2,900	4,000
	7720 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
	7520 Series	2,000	4,000
	7720 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
	7520 Series	1,000	2,000
	7720 Series	1,000	2,000

Number of IS-IS interfaces (NNIs)	Product	I-SIDs with vIST configured on the platform	I-SIDs without vIST configured on the platform
72	5320 Series	n/a	n/a
	5420 Series	n/a	n/a
	5520 Series	750	1,500
	5720 Series	750	1,500
	7520 Series	750	1,500
	7720 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
	7520 Series	550	1,100
	7720 Series	550	1,100
128	5320 Series	n/a	n/a
	5420 Series	n/a	n/a
	5520 Series	450	900
	5720 Series	450	900
	7520 Series	450	900
	7720 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
	7520 Series	n/a	n/a
	7720 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 by means of 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.

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 133

[Compatibility with ExtremeCloud IQ Site Engine](#) on page 133

[Feature-Based Licensing](#) on page 133

[Memory Usage](#) on page 134

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.

For the most current information on switches supported by ExtremeCloud IQ, see [ExtremeCloud™ IQ Release Notes](#).

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](#).

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 switches 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.

For more information about licensing including feature inclusion, order codes, and how to load a license file, see [Fabric Engine User Guide](#).

Memory Usage

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



Known Issues and Restrictions

[Known Issues for this Release](#) on page 135

[Restrictions and Expected Behaviors](#) on page 156

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

Known Issues for this Release

This section identifies the known issues in this release.

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 webserver 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-1280	The following error message occurs when performing shutdown/no-shutdown commands continuously: IO1 [05/02/14 06:59:55.178:UTC] 0x0011c525 00000000 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.

Issue number	Description	Workaround
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-1335	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: <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.

Issue number	Description	Workaround
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: Only 24 Layer 3 VSNs can be configured.	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 .Ip 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 .Ip maps when using Fabric Extend. With default ingress and egress .Ip 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 .Ip 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.

Issue number	Description	Workaround
VOSS-1584	The show debug-file all command is missing.	None.
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:FAIL 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:fe55:5088 cid 2361 owner OSPF </pre>	Disable the alternate path.
VOSS-2117	<p>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.</p>	Disable and re-enable IGMP Snooping on the interface.
VOSS-2128	<p>EAP Security and Authentication EDM tabs display additional information with internal values populated, which is not useful for the end user.</p>	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	<p>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</p>	None.
VOSS-2208	<p>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.</p>	None.

Issue number	Description	Workaround
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-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-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-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.
VOSS-4728	If you remove and recreate an IS-IS instance on an NNI port with auto-negotiation enabled in addition to vIST 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.

Issue number	Description	Workaround
VOSS-5130	Disabling and immediately enabling IS-IS results in the following log message: <code>PLSBFIB ERROR: /vob/cb/nd_protocols/plsb/lib/plsbFib.cpp(line 1558) unregisterLocalInfo() local entry does not exist. key(0xfda010000ffffa40)</code>	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.
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.

Issue number	Description	Workaround
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-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.
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 SSL negotiation with client successful messages during your EDM session. The system displays this message, each time a successful SSL_Handshake 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-9921	Bootup redirection timeout is longer than the UNI port (SMLT) unlock timer. If both v1ST nodes boot together in factory default configuration fabric mode or without a nickname, the v1ST ports will not enable for up to 4 minutes. During the delay the nickname server is unreachable and v1ST is not online.	None.

Issue number	Description	Workaround
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 MHSAs 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-10574	IS-IS sys-name output is not truncated for show isis spbm nick-name or show ip route commands. If a long character sys-name is in use, the full sys-name display can cause misalignment of the output columns.	None.
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-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.

Issue number	Description	Workaround
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 .
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 ixgbevf 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-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-13794	You cannot use SFTP to transfer files larger than 2 GB to the switch.	Use SCP.
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.

Issue number	Description	Workaround
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-14597	Ping (originated from local CP) fails for jumbo frames on Layer 3 VSN interface.	None.
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.
VOSS-15391	An SNMP walk on the rcIcmpSnoopTraceTable table will fail with an OID not increasing error. CLI and EDM are unaffected by this issue.	None.
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.

Issue number	Description	Workaround
VOSS-15812	Layer 3 VSN IPv4 BGP (and static) routes having their next-hops resolved using IS-IS routes could result in traffic loss.	<p>Choose the following workarounds, based on your deployment and needs:</p> <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 workaround scenarios are suitable for your deployment, do not use internal Border Gateway Protocol (iBGP) peering.
VOSS-15878	VSP 4900 Series and 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-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-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.

Issue number	Description	Workaround
VOSS-18023	<p>The management port on the 5520 switch does not support Auto-MDIX (the automatic detection of transmit and received twisted pairs).</p> <p>As a best practice, enable the default auto-negotiation setting on the management port.</p> <p>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	<p>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.</p>	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.</p> <p>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.
VOSS-18360	<p>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.</p>	None.

Issue number	Description	Workaround
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-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 Brouter of the VIST peer.

Issue number	Description	Workaround
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.
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, and VSP 7400 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.

Issue number	Description	Workaround
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-25023	5520 Series, 5420 Series, and 5320 Series platforms can reach 100% CPU utilization during inband transfer (FTP, SFTP, and SCP).	None.
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-25288	Secure boot information for 5720 Series, 7520 Series, and 7720 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 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-26099	MACsec Key Agreement (MKA) MACsec does not operate properly when you enable and disable MKA MACsec on the port 15-20 times.	None.

Issue number	Description	Workaround
VOSS-26122	Intermittently, some CLI commands related to sFlow functionality do not display in the CLI log.	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-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.
VOSS-27643	On 5320 Series, packet port statistics do not increment for multicast traffic ingressing Layer 3 Fabric Extend NNI.	As a workaround, calculate the number of packets from the total number of bytes received.
VOSS-27784	Layer 3 VSN traffic continues to flow after you delete IP addresses in dual stack scenarios.	None.
VOSS-27875	On 7520-48XT-6C copper ports(1/1-1/48) with SLPP enabled, the port LED state is off.	None.

Issue number	Description	Workaround
VOSS-28101	<p>The loss of IP BGP in-route-map and out-route-map from config when you upgrade to Release 8.5.x or later is due to the removal of the following legacy commands in Release 8.5.x that were not needed on newer platforms:</p> <ul style="list-style-type: none"> • ip bgp out-route-map • ip bgp out-route-map 	As a workaround, apply incoming and outgoing route-maps for BGB peers or peer groups.
VOSS-28437	Layer 3 routed traffic is discarded in a square topology with two pairs of vIST DVR controllers in different domains when traffic should reach the diagonal switch.	As a workaround, save the configuration file with the NNI-MSTP flag configured and reboot the system.
VOSS-28241	For a routed Gigabit Ethernet interface, traffic doubles on vIST peers if you issue the action flushALL command.	None.
VOSS-28525	DHCP clients fail to receive an IP address in scenarios with VRRP over SMLT when SMLT goes down and the DHCP interface is configured to broadcast.	As a workaround, disable broadcast on the DHCP relay.
VOSS-28625	<p>Boundary Nodes return VRRP packets into the originating area and cause warning messages to display. The issue occurs if you create the following ACL rule on a Multi-area SPB Boundary Node:</p> <pre style="background-color: #f0f0f0; padding: 5px;"> filter acl 1 type inVsn matchType both filter acl i-sid 1 12990020 filter acl ace 1 1 filter acl ace action 1 1 permit monitor-isid-offset 1 filter acl ace ethernet 1 1 ether-type eq ip filter acl ace 1 1 enable </pre> <p>The issue is caused by the interoperability of this specific ACL configured to mirror the I-SID traffic, and the Multi-area filters.</p>	<p>Remove the ACL used to mirror I-SID traffic on the boundary node. Use Fabric RSPAN (Mirror to I-SID) to achieve similar functionality.</p> <p>Alternatively, use matchtype "uniOnly" instead of "both".</p>
VOSS-28672	IPFIX does not learn MCoSPB NNI-UNI flows on 7520 Series, 7720 Series, and VSP 7400 Series.	None.

Issue number	Description	Workaround
VOSS-29287	Interoperability issues can occur between VOSS/Fabric Engine switches and ExtremeXOS/Switch Engine switches when you use MACsec MKA and disable SCI tagging on both ends. Disabling SCI tagging on both ends works for ExtremeXOS/Switch Engine if the VOSS/Fabric Engine version is earlier than 8.7.	None.
VOSS-29711	If you enter a delayed reboot command for a device with at least one active RADIUS Accounting session, the switch does not send the RADIUS Accounting Stop or RADIUS Accounting Off packets, and console traces display on the screen.	None.
VOSS-29799	Using ZTP+ onboarding with the Management Interface value configured as Management Service for a C-VLAN service does not work. The C-VLAN is created but the management port does not move to the C-VLAN.	Onboarding with Management Service for a DvR Leaf is limited to S-UNI services; you cannot use C-VLAN for a DvR Leaf. For non-DvR and DvR Controllers, change the I-SID after onboarding.
VOSS-30117	On 5520 ACDC models, the XN-DCPWR-550W-BF and XN-DCPWR-550W-FB power supplies do not properly report voltage and amperage values.	None.
VOSS-30195	A potential LLDP flood issue can occur with certain third-party unmanaged devices on Auto-sense ports.	Eliminate the cause of flooding.
VOSS-30222	SSH connection is currently unavailable through Layer 2 FE Tunnel or Layer 3 FE Tunnel on the 5320 Series and 5420 Series.	Enable IPv6 Shortcuts.
VOSS-30292	If IPv6 Shortcuts are explicitly disabled, SSH connections will not work on VSP 4900 Series.	Enable IPv6 Shortcuts.
VOSS-30296	You cannot use SNMP to configure a RADIUS server FQDN with more than 113 characters.	Use CLI or EDM to configure the FQDN.
VOSS-30864	After the switch boots, for a short period of time, some IP Shortcut and IP VPN routes may not be installed if the IP Shortcut or IP VPN restart is not immediately followed by an IS-IS computation. This situation is temporary. After the next IS-IS computation, whether triggered or periodic, all routes are installed in the RTM as expected.	<p>If the issue occurs, you can:</p> <ul style="list-style-type: none"> • Wait for the IS-IS computation to be triggered, with a maximum waiting period of 900 seconds. <p style="text-align: center;">OR</p> <ul style="list-style-type: none"> • Disable and reenable IS-IS. <p>To avoid the issue, configure an IP source address for IP Shortcuts.</p>

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 48: 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.

Table 48: General restrictions (continued)

Issue number	Description	Workaround
VOSS-2166	The IPsec security association (SA) configuration has a NULL Encryption option under the Encrypt-algo parameter. Currently, you must fill the encryptKey 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.
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	OVSDDB 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.

Table 48: General restrictions (continued)

Issue number	Description	Workaround
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.
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	<p>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.</p> <p>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 .</p>	If additional performance is desired, upgrade the VM guest OS with an Ethernet device driver that supports 10 Gbps interfaces.
VOSS-12151	<p>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.</p> <p>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.</p>	After you connect the VM to the software VTEP, the issue is not seen.
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.

Table 48: General restrictions (continued)

Issue number	Description	Workaround
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
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
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

Table 48: General restrictions (continued)

Issue number	Description	Workaround
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 48: 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>
wi01171670	Telnet packets get encrypted on MACsec-enabled ports.	None.
wi01210217	The command show eapol auth-stats displays LAST-SRC-MAC for NEAP sessions incorrectly.	n/a
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
wi01213066 wi01213374	EAP and NEAP are not supported on brouter ports.	n/a

Table 48: General restrictions (continued)

Issue number	Description	Workaround
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.
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.

Table 48: General restrictions (continued)

Issue number	Description	Workaround
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.
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 49: 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 49: 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 50: 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:

- Fabric Engine 8.10.4

Issue number	Description
CFD-10229	Dropping ARP reply packet destined for its peer when ingress in different VLAN and needs to be bridged out to the destined VLAN
CFD-10804	Only 5320-48P-8XE and 5320-48T-8XE support more than one VRF with IP configuration (either GRT or user-created). Because of this restriction, Auto-sense cannot create the automatic SD-WAN VRF configuration for the 16- or 24-port models if an IP configuration already exists. For configuration information, see Specify the VRF for Auto-sense ExtremeCloud SD-WAN Configuration on page 40.
CFD-10917	FDB entries not flushed when TCN received on a ring port
CFD-11001	A static S-UNI cannot be added to an I-SID mapped to a dynamic VLAN
CFD-11062	Error in the console <code>Error parsing '/intflash/khi/khi_boot_count'!</code> .
CFD-11178	The <code>show sys-info fan</code> command displays information intermittently.
CFD-11280	VSP 7400 Series: Extreme Optics reporting 70+°C and Fan speed remains low.
VOSS-29220	In a scaled Multi-area SPB topology, after an event like an NNI link down, the fail over time for multicast traffic can take up to 23-25 seconds when using 5720 Series switches as the boundary node pair.



Related Information

[MIB Changes](#) on page 166

MIB Changes

Deprecated MIBs

Table 51: 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
rcDvrBackboneEntriesArea	1.3.6.1.4.1.2272.1.219.8.1.12	9.0
rcDvrBackboneMemberArea	1.3.6.1.4.1.2272.1.219.9.1.6	9.0
rcDvrBackboneMultiAreaVnodeEntriesArea	1.3.6.1.4.1.2272.1.219.10.1.12	9.0

Modified MIBs

Table 52: Common

Object Name	Object OID	Modified in Release	Modification
SnpxChassisType		9.0	ADD ENUM: m552024TACDC, m552048TACDC, m552024XACDC, m552048SEACDC
avFabricAttachElementType	1.3.6.1.4.1.45.5.46.1.2	9.0	ADD_ENUM: faRing(18)
avFabricAttachDiscElemsElementType	1.3.6.1.4.1.45.5.46.1.11.1.2	9.0	ADD_ENUM: faRing(18)

Table 52: Common (continued)

Object Name	Object OID	Modified in Release	Modification
rcSysActionL1	1.3.6.1.4.1.2272.1.1.86	9.0	OTHER: Update description for revokeLicense10G4P, revokeLicense10G8P, not supported starting with release 9.0
rcSysActionL1	1.3.6.1.4.1.2272.1.1.86	9.0	ADD ENUM:revokeLicensePremier, revokeLicenseMacsec for 7x20
rcSysActionRwa	1.3.6.1.4.1.2272.1.1.89	9.0	OTHER: ADD ENUM: softResetDelay, softResetCancel
rcChasType	1.3.6.1.4.1.2272.1.4.1	9.0	ADD ENUM: a552024TACDC, a552048TACDC, a552048SEACDC, a552024XACDC, a752048YE8CE
rcPortAutoSenseState	1.3.6.1.4.1.2272.1.4.10.1.1.132	9.0	ADD ENUM: nniPending(13), sdWan(14), sdWanPending(15)
rcPortAutoSenseState	1.3.6.1.4.1.2272.1.4.10.1.1.134	9.0	ADD_ENUM: faRing(16)
rcIsgLogicalInterfaceSrcIPAddr	1.3.6.1.4.1.2272.1.63.26.1.31	9.0	OTHER: Updated description to be available on all platforms
rc2kBootConfigEnableFactoryDefaultsMode	1.3.6.1.4.1.2272.1.100.5.1.60	9.0	ADD_NEW_VALUES: Add value zero-touch-config-only to factorydefaults options
rc2kCardFrontType	1.3.6.1.4.1.2272.1.100.6.1.2	9.0	ADD ENUM: fabricEngine5520x24TACDC, fabricEngine5520x48TACDC, fabricEngine5520x48SEACDC, fabricEngine5520x24XACDC, fabricEngine752048YE8CE
rcVossSystemMgmtPortLedStatus	1.3.6.1.4.1.2272.1.101.1.1.1.1	9.0	OTHER: Update description to include 7520-48YE-8CE

Table 52: Common (continued)

Object Name	Object OID	Modified in Release	Modification
rcVlanMvvpnIsidStatus	1.3.6.1.4.1.2272.1.3.2.1.84	9.0.2	ADD_NEW_VALUE: not-configured(3) OTHER: Updated description
rcMACSecConnectivityAssociationName	1.3.6.1.4.1.2272.1.88.1.1.2	9.0.2	CHANGE_RANGE: Changed the range from 5..16 to 5..32
rcMACSecIfCAName	1.3.6.1.4.1.2272.1.88.2.1.1	9.0.2	CHANGE_RANGE: Changed the range from 5..16 to 5..32
rcIpAdEntIfType	1.3.6.1.4.1.2272.1.8.2.1.10	9.0.3	CHANGE: index MAX-ACCESS level: from read-only to read-write rcIpAdEntIfType 1.3.6.1.4.1.2272.1.8.2.1.10 OTHER: Update description to include the new values added to enum

Table 53: 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
rcIshPlsbSmltBmac	1.3.6.1.4.1.2272.1.63.4.1.10	8.6.1	Not supported on 5320 platform
rcIshPlsbSmltPeerSysId	1.3.6.1.4.1.2272.1.63.4.1.11	8.6.1	Not supported on 5320 platform
rcIpConfRsmItEnable	1.3.6.1.4.1.2272.1.8.1.1.21	8.6.1	Not supported on 5320 platform
rcIpConfRsmItTable	1.3.6.1.4.1.2272.1.8.1.11	8.6.1	Not supported on 5320 platform
rcIpRsmItEdgeSupportEnable	1.3.6.1.4.1.2272.1.8.26.1.2	8.6.1	Not supported on 5320 platform

Table 53: 5320 Series (continued)

Object Name	Object OID	Modified in Release	Modification
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 54: 5420 Series

Object Name	Object OID	Modified in Release	Modification
SnpxChassisType		8.6	<p>ADD ENUM: m532048T8XE, m532048P8XE, m532024T8XE, m532024P8XE, m532016P4XE, m532016P4XEDC</p> <p>OTHER: Replace "Virtual Services" with "Extreme Networks Fabric Engine" and "VOSS" with "FabricEngine" in comments only for Universal Hardware</p> <p>OTHER: Rebranding for Universal Hardware: Change enum values from m552048TVOSS, m552048WVOSS, m552012MW36WVOSS, m552024TVOSS, m552024WVOSS, m552024XVOSS, m552048SEVOSS to m552048T, m552048W, m552012MW36W, m552024T, m552024W,</p>

Table 54: 5420 Series (continued)

Object Name	Object OID	Modified in Release	Modification
			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, a532016P4XEDCVOS S
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
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 54: 5420 Series (continued)

Object Name	Object OID	Modified in Release	Modification
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
bspePethPsePortPowerClassifications	1.3.6.1.4.1.45.5.8.1.1.15	8.10	OTHER: Updated description to include 5720 platform

Table 55: 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 55: 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
bspePethPsePortPowerClassifications	1.3.6.1.4.1.45.5.8.1.1.15	8.10	OTHER: Updated description to include 5720 platform

Table 56: 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)

Table 56: 5720 Series (continued)

Object Name	Object OID	Modified in Release	Modification
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
bspePethPsePortPowerClassifications	1.3.6.1.4.1.45.5.8.1.1.15	8.10	OTHER: Updated description to include 5720 platform

Table 57: 7520 Series

Object Name	Object OID	Modified in Release	Modification
rcChasType	1.3.6.1.4.1.2272.1.4.1	8.10	OTHER: Replace "VOSS" with "FabricEngine" in 7520 and 7720 models values
rc2kCardFrontType	1.3.6.1.4.1.2272.1.100.6.1.2	8.10	OTHER: Replace "VOSS" with "FabricEngine" in 7520 and 7720 models values
rcEapMultiHostStatusDynamicSettings	1.3.6.1.4.1.2272.1.57.4.1.11	8.10	CHANGE_RANGE: Changed range from 0..100 to 0..150
rcEapPortDynamicSettings	1.3.6.1.4.1.2272.1.57.6.1.12	8.10	CHANGE_RANGE: Changed range from 0..100 to 0..150
rcVossSystemCardLedStatus	1.3.6.1.4.1.2272.1.101.1.1.5.1.4	8.10	ADD_NEW_VALUE: blueSteady(9)
rcVossSystemMgmtPortLedStatus	1.3.6.1.4.1.2272.1.101.1.1.1.1	8.10	OTHER: Update description to include 7520 and 7720
rc2kChassisPortLed3Status	1.3.6.1.4.1.2272.1.100.1.14	8.10	OTHER: Update 7520 and 7720 description
rc2kBootConfigAdvancedFeatureBwReservation	1.3.6.1.4.1.2272.1.100.5.1.51	8.10	OTHER: Update description and synchronize between files

Table 58: 7720 Series

Object Name	Object OID	Modified in Release	Modification
rcChasType	1.3.6.1.4.1.2272.1.4.1	8.10	OTHER: Replace "VOSS" with "FabricEngine" in

Table 58: 7720 Series (continued)

Object Name	Object OID	Modified in Release	Modification
			7520 and 7720 models values
rc2kCardFrontType	1.3.6.1.4.1.2272.1.100.6.1.2	8.10	OTHER: Replace "VOSS" with "FabricEngine" in 7520 and 7720 models values
rcEapMultiHostStatusDynamicSettings	1.3.6.1.4.1.2272.1.57.4.1.11	8.10	CHANGE_RANGE: Changed range from 0..100 to 0..150
rcEapPortDynamicSettings	1.3.6.1.4.1.2272.1.57.6.1.12	8.10	CHANGE_RANGE: Changed range from 0..100 to 0..150
rcVossSystemCardLedStatus	1.3.6.1.4.1.2272.1.101.1.1.5.1.4	8.10	ADD_NEW_VALUE: blueSteady(9)
rcVossSystemMgmtPortLedStatus	1.3.6.1.4.1.2272.1.101.1.1.1.1	8.10	OTHER: Update description to include 7520 and 7720
rc2kChassisPortLed3Status	1.3.6.1.4.1.2272.1.100.1.14	8.10	OTHER: Update 7520 and 7720 description
rc2kBootConfigAdvancedFeatureBwReservation	1.3.6.1.4.1.2272.1.100.5.1.51	8.10	OTHER: Update description and synchronize between files

New MIBs

Table 59: Common

Object Name	Object OID	New in VOSS Release
avFabricAttachPortTCNEnable	1.3.6.1.4.1.45.5.46.1.6.1.8	9.0
rcDhcpServer	1.3.6.1.4.1.2272.1.232	9.0
rcDhcpServerMib	1.3.6.1.4.1.2272.1.232.1	9.0
rcDhcpServerNotifications	1.3.6.1.4.1.2272.1.232.1.0	9.0
rcDhcpServerObjects	1.3.6.1.4.1.2272.1.232.1.1	9.0
rcDhcpServerGlobal	1.3.6.1.4.1.2272.1.232.1.1.1	9.0
rcDhcpServerSubnetTable	1.3.6.1.4.1.2272.1.232.1.1.2	9.0
rcDhcpServerHostTable	1.3.6.1.4.1.2272.1.232.1.1.3	9.0
rcDhcpServerGlobalDnsTable	1.3.6.1.4.1.2272.1.232.1.1.4	9.0
rcDhcpServerGlobalNtpTable	1.3.6.1.4.1.2272.1.232.1.1.5	9.0
rcDhcpServerSubnetRouterTable	1.3.6.1.4.1.2272.1.232.1.1.6	9.0

Table 59: Common (continued)

Object Name	Object OID	New in VOSS Release
rcDhcpServerSubnetDnsTable	1.3.6.1.4.1.2272.1.232.1.1.7	9.0
rcDhcpServerSubnetNtpTable	1.3.6.1.4.1.2272.1.232.1.1.8	9.0
rcDhcpServerGlobalCustomOptionDefTable	1.3.6.1.4.1.2272.1.232.1.1.9	9.0
rcDhcpServerGlobalCustomOptionDataTable	1.3.6.1.4.1.2272.1.232.1.1.10	9.0
rcDhcpServerSubnetCustomOptionDataTable	1.3.6.1.4.1.2272.1.232.1.1.11	9.0
rcWebSSLRenegotiation	1.3.6.1.4.1.2272.1.18.38	9.0
rcSysResetDelayTimeout	10.101.18.21 1.3.6.1.4.1.2272.1.1.130	9.0
rcEapMultiHostStatusMacClear	1.3.6.1.4.1.2272.1.57.4.1.14	9.0.2
rcAutoSenseFaProxyRingMgmtIsid	1.3.6.1.4.1.2272.1.231.1.1.1.29	9.0.2
rcAutoSenseFaProxyRingMgmtCvid	1.3.6.1.4.1.2272.1.231.1.1.1.30	9.0.2
rcDhcpServerGlobalVendorOptionDefTable	1.3.6.1.4.1.2272.1.232.1.1.12	9.0.2
rcDhcpServerGlobalVendorOptionDataTable	1.3.6.1.4.1.2272.1.232.1.1.13	9.0.2
rcDhcpServerVendorClassTable	1.3.6.1.4.1.2272.1.232.1.1.15	9.0.2
rcDhcpServerVendorClassCustomOptionDataTable	1.3.6.1.4.1.2272.1.232.1.1.16	9.0.2
rcDhcpServerVendorClassVendorOptionDataTable	1.3.6.1.4.1.2272.1.232.1.1.17	9.0.2
rcLsisLogicalInterfaceMAVirtualLink	1.3.6.1.4.1.2272.1.63.26.1.35	9.0.3
rcLldpXMedLocMediaPolicyTable	1.3.6.1.4.1.2272.1.220.1.2.5	9.0.3
rcLldpXMedLocMediaPolicyLocalPortNum	1.3.6.1.4.1.2272.1.220.1.2.5.1.1	9.0.3
rcLldpXMedLocMediaPolicyAppType	1.3.6.1.4.1.2272.1.220.1.2.5.1.2	9.0.3
rcLldpXMedLocMediaPolicyVlanID	1.3.6.1.4.1.2272.1.220.1.2.5.1.3	9.0.3
rcLldpXMedLocMediaPolicyPriority	1.3.6.1.4.1.2272.1.220.1.2.5.1.4	9.0.3
rcLldpXMedLocMediaPolicyDscp	1.3.6.1.4.1.2272.1.220.1.2.5.1.5	9.0.3

Table 59: Common (continued)

Object Name	Object OID	New in VOSS Release
rcLldpXMedLocMediaPolicyRowStatus	1.3.6.1.4.1.2272.1.220.1.2.5.1.6	9.0.3
rcLldpXMedLocMediaPolicyTagged	1.3.6.1.4.1.2272.1.220.1.2.5.1.7	9.0.3

Table 60: 5320 Series

Object Name	Object OID	New in Release
rcVlanMvpngIsidOffset	1.3.6.1.4.1.2272.1.3.2.1.86	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
rcDiagVctTable	1.3.6.1.4.1.2272.1.23.4	8.10
rcDiagVctEntry	1.3.6.1.4.1.2272.1.23.4.1	8.10
rcDiagVctIfIndex	1.3.6.1.4.1.2272.1.23.4.1.1	8.10
rcDiagVctNormalCableLength	1.3.6.1.4.1.2272.1.23.4.1.2	8.10
rcDiagVctCableStatus	1.3.6.1.4.1.2272.1.23.4.1.4	8.10
rcDiagVctPair1Status	1.3.6.1.4.1.2272.1.23.4.1.5	8.10
rcDiagVctPair1ErrLength	1.3.6.1.4.1.2272.1.23.4.1.6	8.10
rcDiagVctPair2Status	1.3.6.1.4.1.2272.1.23.4.1.7	8.10
rcDiagVctPair2ErrLength	1.3.6.1.4.1.2272.1.23.4.1.8	8.10
rcDiagVctPair3Status	1.3.6.1.4.1.2272.1.23.4.1.9	8.10
rcDiagVctPair3ErrLength	1.3.6.1.4.1.2272.1.23.4.1.10	8.10
rcDiagVctPair4Status	1.3.6.1.4.1.2272.1.23.4.1.11	8.10
rcDiagVctPair4ErrLength	1.3.6.1.4.1.2272.1.23.4.1.12	8.10
rcDiagVctStartTest	1.3.6.1.4.1.2272.1.23.4.1.13	8.10
rcDiagVctTestDone	1.3.6.1.4.1.2272.1.23.4.1.14	8.10
rcDiagVctCableLength	1.3.6.1.4.1.2272.1.23.4.1.16	8.10
rcIshisPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcIshisPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcIshisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	8.10

Table 60: 5320 Series (continued)

Object Name	Object OID	New in Release
rcIsisLogicalInterfacelsisMtu	1.3.6.1.4.1.2272.1.63.26.1.33	8.10
rcLldpXMedLocMediaPolicyExtendedTable	1.3.6.1.4.1.2272.1.220.1.2.4	8.10
rcLldpXMedLocMediaPolicyExtendedEntry	1.3.6.1.4.1.2272.1.220.1.2.4.1	8.10
rcLldpXMedLocMediaPolicyExtendedOrigin	1.3.6.1.4.1.2272.1.220.1.2.4.1.1	8.10
rcAutoSenseSdWanVrfName	1.3.6.1.4.1.2272.1.231.1.1.1.32	9.0.3

Table 61: 5420 Series

Object Name	Object OID	New in Release
rcVlanMvvpnIsidOffset	1.3.6.1.4.1.2272.1.3.2.1.86	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
rcDiagVctTable	1.3.6.1.4.1.2272.1.23.4	8.10
rcDiagVctEntry	1.3.6.1.4.1.2272.1.23.4.1	8.10
rcDiagVctIfIndex	1.3.6.1.4.1.2272.1.23.4.1.1	8.10
rcDiagVctNormalCableLength	1.3.6.1.4.1.2272.1.23.4.1.2	8.10
rcDiagVctCableStatus	1.3.6.1.4.1.2272.1.23.4.1.4	8.10
rcDiagVctPair1Status	1.3.6.1.4.1.2272.1.23.4.1.5	8.10
rcDiagVctPair1ErrLength	1.3.6.1.4.1.2272.1.23.4.1.6	8.10
rcDiagVctPair2Status	1.3.6.1.4.1.2272.1.23.4.1.7	8.10
rcDiagVctPair2ErrLength	1.3.6.1.4.1.2272.1.23.4.1.8	8.10
rcDiagVctPair3Status	1.3.6.1.4.1.2272.1.23.4.1.9	8.10
rcDiagVctPair3ErrLength	1.3.6.1.4.1.2272.1.23.4.1.10	8.10
rcDiagVctPair4Status	1.3.6.1.4.1.2272.1.23.4.1.11	8.10
rcDiagVctPair4ErrLength	1.3.6.1.4.1.2272.1.23.4.1.12	8.10
rcDiagVctStartTest	1.3.6.1.4.1.2272.1.23.4.1.13	8.10
rcDiagVctTestDone	1.3.6.1.4.1.2272.1.23.4.1.14	8.10
rcDiagVctCableLength	1.3.6.1.4.1.2272.1.23.4.1.16	8.10

Table 61: 5420 Series (continued)

Object Name	Object OID	New in Release
rcIsisPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcIsisPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcIsisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	8.10
rcIsisLogicalInterfacelsisMtu	1.3.6.1.4.1.2272.1.63.26.1.33	8.10
rcLldpXMedLocMediaPolicyExtendedTable	1.3.6.1.4.1.2272.1.220.1.2.4	8.10
rcLldpXMedLocMediaPolicyExtendedEntry	1.3.6.1.4.1.2272.1.220.1.2.4.1	8.10
rcLldpXMedLocMediaPolicyExtendedOrigin	1.3.6.1.4.1.2272.1.220.1.2.4.1.1	8.10
rcChasPowerSupplyDetailVoltageIn	1.3.6.1.4.1.2272.1.4.8.2.1.16	9.0.2
rcChasPowerSupplyDetailVoltageOut	1.3.6.1.4.1.2272.1.4.8.2.1.17	9.0.2
rcChasPowerSupplyDetailCurrentIn	1.3.6.1.4.1.2272.1.4.8.2.1.18	9.0.2
rcChasPowerSupplyDetailCurrentOut	1.3.6.1.4.1.2272.1.4.8.2.1.19	9.0.2
rcChasPowerSupplyDetailPowerIn	1.3.6.1.4.1.2272.1.4.8.2.1.20	9.0.2
rcChasPowerSupplyDetailPowerOut	1.3.6.1.4.1.2272.1.4.8.2.1.21	9.0.2

Table 62: 5520 Series

Object Name	Object OID	New in Release
rcDiagVctTable	1.3.6.1.4.1.2272.1.23.4	8.10
rcDiagVctEntry	1.3.6.1.4.1.2272.1.23.4.1	8.10
rcDiagVctIfIndex	1.3.6.1.4.1.2272.1.23.4.1.1	8.10
rcDiagVctNormalCableLength	1.3.6.1.4.1.2272.1.23.4.1.2	8.10
rcDiagVctCableStatus	1.3.6.1.4.1.2272.1.23.4.1.4	8.10
rcDiagVctPair1Status	1.3.6.1.4.1.2272.1.23.4.1.5	8.10
rcDiagVctPair1ErrLength	1.3.6.1.4.1.2272.1.23.4.1.6	8.10
rcDiagVctPair2Status	1.3.6.1.4.1.2272.1.23.4.1.7	8.10
rcDiagVctPair2ErrLength	1.3.6.1.4.1.2272.1.23.4.1.8	8.10
rcDiagVctPair3Status	1.3.6.1.4.1.2272.1.23.4.1.9	8.10
rcDiagVctPair3ErrLength	1.3.6.1.4.1.2272.1.23.4.1.10	8.10
rcDiagVctPair4Status	1.3.6.1.4.1.2272.1.23.4.1.11	8.10

Table 62: 5520 Series (continued)

Object Name	Object OID	New in Release
rcDiagVctPair4ErrLength	1.3.6.1.4.1.2272.1.23.4.1.12	8.10
rcDiagVctStartTest	1.3.6.1.4.1.2272.1.23.4.1.13	8.10
rcDiagVctTestDone	1.3.6.1.4.1.2272.1.23.4.1.14	8.10
rcDiagVctCableLength	1.3.6.1.4.1.2272.1.23.4.1.16	8.10
rcLsisPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcLsisPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcLsisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	8.10
rcLsisLogicalInterfacelSisMtu	1.3.6.1.4.1.2272.1.63.26.1.33	8.10
rcLldpXMedLocMediaPolicyExtendedTable	1.3.6.1.4.1.2272.1.220.1.2.4	8.10
rcLldpXMedLocMediaPolicyExtendedEntry	1.3.6.1.4.1.2272.1.220.1.2.4.1	8.10
rcLldpXMedLocMediaPolicyExtendedOrigin	1.3.6.1.4.1.2272.1.220.1.2.4.1.1	8.10
rcChasPowerSupplyDetailVoltageIn	1.3.6.1.4.1.2272.1.4.8.2.1.16	9.0.2
rcChasPowerSupplyDetailVoltageOut	1.3.6.1.4.1.2272.1.4.8.2.1.17	9.0.2
rcChasPowerSupplyDetailCurrentIn	1.3.6.1.4.1.2272.1.4.8.2.1.18	9.0.2
rcChasPowerSupplyDetailCurrentOut	1.3.6.1.4.1.2272.1.4.8.2.1.19	9.0.2
rcChasPowerSupplyDetailPowerIn	1.3.6.1.4.1.2272.1.4.8.2.1.20	9.0.2
rcChasPowerSupplyDetailPowerOut	1.3.6.1.4.1.2272.1.4.8.2.1.21	9.0.2
rcAutoSenseSdWanArea	1.3.6.1.4.1.2272.1.231.1.1.1.31	9.0.3
rcAutoSenseSdWanInterfaceTable	1.3.6.1.4.1.2272.1.231.1.1.2	9.0.3
rcAutoSenseSdWanInterfaceIpl	1.3.6.1.4.1.2272.1.231.1.1.2.1.1	9.0.3
rcAutoSenseSdWanInterfaceRowStatus	1.3.6.1.4.1.2272.1.231.1.1.2.1.2	9.0.3
rcAutoSenseSdWanInterfaceArea	1.3.6.1.4.1.2272.1.231.1.1.2.1.3	9.0.3

Table 63: 5720 Series

Object Name	Object OID	New in Release
	1.3.6.1.4.1.2272.1.23.4	8.10

Table 63: 5720 Series (continued)

Object Name	Object OID	New in Release
rcDiagVctTable		
rcDiagVctEntry	1.3.6.1.4.1.2272.1.23.4.1	8.10
rcDiagVctIfIndex	1.3.6.1.4.1.2272.1.23.4.1.1	8.10
rcDiagVctNormalCableLength	1.3.6.1.4.1.2272.1.23.4.1.2	8.10
rcDiagVctCableStatus	1.3.6.1.4.1.2272.1.23.4.1.4	8.10
rcDiagVctPair1Status	1.3.6.1.4.1.2272.1.23.4.1.5	8.10
rcDiagVctPair1ErrLength	1.3.6.1.4.1.2272.1.23.4.1.6	8.10
rcDiagVctPair2Status	1.3.6.1.4.1.2272.1.23.4.1.7	8.10
rcDiagVctPair2ErrLength	1.3.6.1.4.1.2272.1.23.4.1.8	8.10
rcDiagVctPair3Status	1.3.6.1.4.1.2272.1.23.4.1.9	8.10
rcDiagVctPair3ErrLength	1.3.6.1.4.1.2272.1.23.4.1.10	8.10
rcDiagVctPair4Status	1.3.6.1.4.1.2272.1.23.4.1.11	8.10
rcDiagVctPair4ErrLength	1.3.6.1.4.1.2272.1.23.4.1.12	8.10
rcDiagVctStartTest	1.3.6.1.4.1.2272.1.23.4.1.13	8.10
rcDiagVctTestDone	1.3.6.1.4.1.2272.1.23.4.1.14	8.10
rcDiagVctCableLength	1.3.6.1.4.1.2272.1.23.4.1.16	8.10
rcLsisPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcLsisPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcLsisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	8.10
rcLsisLogicalInterfacelSisMtu	1.3.6.1.4.1.2272.1.63.26.1.33	8.10
rcLldpXMedLocMediaPolicyExtendedTable	1.3.6.1.4.1.2272.1.220.1.2.4	8.10
rcLldpXMedLocMediaPolicyExtendedEntry	1.3.6.1.4.1.2272.1.220.1.2.4.1	8.10
rcLldpXMedLocMediaPolicyExtendedOrigin	1.3.6.1.4.1.2272.1.220.1.2.4.1.1	8.10
rcChasPowerSupplyDetailVoltageIn	1.3.6.1.4.1.2272.1.4.8.2.1.16	9.0.2
rcChasPowerSupplyDetailVoltageOut	1.3.6.1.4.1.2272.1.4.8.2.1.17	9.0.2
rcChasPowerSupplyDetailCurrentIn	1.3.6.1.4.1.2272.1.4.8.2.1.18	9.0.2
rcChasPowerSupplyDetailCurrentOut	1.3.6.1.4.1.2272.1.4.8.2.1.19	9.0.2
rcChasPowerSupplyDetailPowerIn	1.3.6.1.4.1.2272.1.4.8.2.1.20	9.0.2

Table 63: 5720 Series (continued)

Object Name	Object OID	New in Release
rcChasPowerSupplyDetailPowerOut	1.3.6.1.4.1.2272.1.4.8.2.1.21	9.0.2
rcAutoSenseSdWanArea	1.3.6.1.4.1.2272.1.231.1.1.1.31	9.0.3
rcAutoSenseSdWanInterfaceTable	1.3.6.1.4.1.2272.1.231.1.1.2	9.0.3
rcAutoSenseSdWanInterfaceIpl	1.3.6.1.4.1.2272.1.231.1.1.2.1.1	9.0.3
rcAutoSenseSdWanInterfaceRowStatus	1.3.6.1.4.1.2272.1.231.1.1.2.1.2	9.0.3
rcAutoSenseSdWanInterfaceArea	1.3.6.1.4.1.2272.1.231.1.1.2.1.3	9.0.3

Table 64: 7520 Series

Object Name	Object OID	New in Release
rcDiagVctTable	1.3.6.1.4.1.2272.1.23.4	8.10
rcDiagVctEntry	1.3.6.1.4.1.2272.1.23.4.1	8.10
rcDiagVctIfIndex	1.3.6.1.4.1.2272.1.23.4.1.1	8.10
rcDiagVctNormalCableLength	1.3.6.1.4.1.2272.1.23.4.1.2	8.10
rcDiagVctCableStatus	1.3.6.1.4.1.2272.1.23.4.1.4	8.10
rcDiagVctPair1Status	1.3.6.1.4.1.2272.1.23.4.1.5	8.10
rcDiagVctPair1ErrLength	1.3.6.1.4.1.2272.1.23.4.1.6	8.10
rcDiagVctPair2Status	1.3.6.1.4.1.2272.1.23.4.1.7	8.10
rcDiagVctPair2ErrLength	1.3.6.1.4.1.2272.1.23.4.1.8	8.10
rcDiagVctPair3Status	1.3.6.1.4.1.2272.1.23.4.1.9	8.10
rcDiagVctPair3ErrLength	1.3.6.1.4.1.2272.1.23.4.1.10	8.10
rcDiagVctPair4Status	1.3.6.1.4.1.2272.1.23.4.1.11	8.10
rcDiagVctPair4ErrLength	1.3.6.1.4.1.2272.1.23.4.1.12	8.10
rcDiagVctStartTest	1.3.6.1.4.1.2272.1.23.4.1.13	8.10
rcDiagVctTestDone	1.3.6.1.4.1.2272.1.23.4.1.14	8.10
rcDiagVctCableLength	1.3.6.1.4.1.2272.1.23.4.1.16	8.10
rcIlsisPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcIlsisPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcIlsisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	8.10
rcIlsisLogicalInterfacelIsmtu	1.3.6.1.4.1.2272.1.63.26.1.33	8.10
rcLdpXMedLocMediaPolicyExtendedTable	1.3.6.1.4.1.2272.1.220.1.2.4	8.10

Table 64: 7520 Series (continued)

Object Name	Object OID	New in Release
rcLldpXMedLocMediaPolicyExtendedEntry	1.3.6.1.4.1.2272.1.220.1.2.4.1	8.10
rcLldpXMedLocMediaPolicyExtendedOrigin	1.3.6.1.4.1.2272.1.220.1.2.4.1.1	8.10
rcChasPowerSupplyDetailVoltageIn	1.3.6.1.4.1.2272.1.4.8.2.1.16	9.0.2
rcChasPowerSupplyDetailVoltageOut	1.3.6.1.4.1.2272.1.4.8.2.1.17	9.0.2
rcChasPowerSupplyDetailCurrentIn	1.3.6.1.4.1.2272.1.4.8.2.1.18	9.0.2
rcChasPowerSupplyDetailCurrentOut	1.3.6.1.4.1.2272.1.4.8.2.1.19	9.0.2
rcChasPowerSupplyDetailPowerIn	1.3.6.1.4.1.2272.1.4.8.2.1.20	9.0.2
rcChasPowerSupplyDetailPowerOut	1.3.6.1.4.1.2272.1.4.8.2.1.21	9.0.2
rcAutoSenseSdWanArea	1.3.6.1.4.1.2272.1.231.1.1.1.31	9.0.3
rcAutoSenseSdWanInterfaceTable	1.3.6.1.4.1.2272.1.231.1.1.2	9.0.3
rcAutoSenseSdWanInterfaceIpl	1.3.6.1.4.1.2272.1.231.1.1.2.1.1	9.0.3
rcAutoSenseSdWanInterfaceRowStatus	1.3.6.1.4.1.2272.1.231.1.1.2.1.2	9.0.3
rcAutoSenseSdWanInterfaceArea	1.3.6.1.4.1.2272.1.231.1.1.2.1.3	9.0.3

Table 65: 7720 Series

Object Name	Object OID	New in Release
rcIlsisPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcIlsisPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcIlsisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	8.10
rcIlsisLogicalInterfaceIlsisMtu	1.3.6.1.4.1.2272.1.63.26.1.33	8.10
rcLldpXMedLocMediaPolicyExtendedTable	1.3.6.1.4.1.2272.1.220.1.2.4	8.10
rcLldpXMedLocMediaPolicyExtendedEntry	1.3.6.1.4.1.2272.1.220.1.2.4.1	8.10
rcLldpXMedLocMediaPolicyExtendedOrigin	1.3.6.1.4.1.2272.1.220.1.2.4.1.1	8.10
rcChasPowerSupplyDetailVoltageIn	1.3.6.1.4.1.2272.1.4.8.2.1.16	9.0.2

Table 65: 7720 Series (continued)

Object Name	Object OID	New in Release
rcChasPowerSupplyDetailVoltageOut	1.3.6.1.4.1.2272.1.4.8.2.1.17	9.0.2
rcChasPowerSupplyDetailCurrentIn	1.3.6.1.4.1.2272.1.4.8.2.1.18	9.0.2
rcChasPowerSupplyDetailCurrentOut	1.3.6.1.4.1.2272.1.4.8.2.1.19	9.0.2
rcChasPowerSupplyDetailPowerIn	1.3.6.1.4.1.2272.1.4.8.2.1.20	9.0.2
rcChasPowerSupplyDetailPowerOut	1.3.6.1.4.1.2272.1.4.8.2.1.21	9.0.2
rcAutoSenseSdWanArea	1.3.6.1.4.1.2272.1.231.1.1.1.31	9.0.3
rcAutoSenseSdWanInterfaceTable	1.3.6.1.4.1.2272.1.231.1.1.2	9.0.3
rcAutoSenseSdWanInterfaceIpl	1.3.6.1.4.1.2272.1.231.1.1.2.1.1	9.0.3
rcAutoSenseSdWanInterfaceRowStatus	1.3.6.1.4.1.2272.1.231.1.1.2.1.2	9.0.3
rcAutoSenseSdWanInterfaceArea	1.3.6.1.4.1.2272.1.231.1.1.2.1.3	9.0.3

Obsolete MIBs

Table 66: Common

Object Name	Object OID	Obsolete in Release
rcIpbGpTmpEstablishedNotification	1.3.6.1.4.1.2272.1.8.101.17.0.1	8.10.1
rcIpbGpTmpBackwardTransNotification	1.3.6.1.4.1.2272.1.8.101.17.0.2	8.10.1