



VOSS Release Notes

For VOSS Release 9.0.2

9037801-01 Rev AC
April 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
Channelization and MLT Support with 100G-LR4-QSFP10KM	13
DHCP Server Enhancement.....	13
General Enhancements.....	14
Multicast Lite RADIUS VSA Enhancement	14
Power Reporting Enhancement.....	15
RADIUS Configuration using Fully Qualified Domain Name (FQDN).....	15
Other Changes.....	15
Deprecated Command.....	15
Filenames for this Release.....	15
Upgrade and Downgrade Considerations.....	17
Validated Upgrade Paths.....	17
Switches That Will Not Use Zero Touch Deployment.....	18
Switches That Will Use Zero Touch Deployment	18
Compatible Fabric IPsec Gateway Versions.....	20
Downgrade Considerations.....	20
ExtremeCloud IQ Agent.....	20
Downgrade ExtremeCloud IQ Managed Switches to 9.0.0.0.....	21
Migration to Segmented Management Instance.....	22
Segmented Management Instance Migration and DVR	25
Post Upgrade Configuration for Zero Touch Fabric Configuration and Nickname Assignment.....	25
Network Requirements.....	26
Zero Touch Fabric Configuration Switch.....	27
Hardware and Software Compatibility.....	29
VSP 4450 Series Hardware.....	29
VSP 4900 Series Hardware.....	29

VSP 4900 Series Operational Notes.....	30
Versatile Interface Module Operational Notes.....	31
VIM5-2Y and VIM5-4Y Operational Notes.....	32
VSP 7200 Series Hardware.....	32
VSP 7400 Series Hardware.....	33
VSP 7400 Series Operational Notes.....	33
VSP 8000 Series Hardware.....	34
XA1400 Series Hardware.....	34
Transceivers.....	34
Auto-Negotiation.....	35
Forward Error Correction (FEC).....	35
Power Supply Compatibility.....	35
Scaling.....	36
Layer 2.....	37
Maximum Number of Directed Broadcast Interfaces.....	39
Maximum Number of Microsoft NLB Cluster IP Interfaces.....	39
IP Unicast.....	40
IP Interface Maximums for VSP 4900 Series	43
IP Interface Maximums for VSP 7400 Series.....	43
Layer 3 Route Table Size.....	44
Route Scaling.....	44
IP Multicast.....	45
Distributed Virtual Routing (DvR).....	47
VXLAN Gateway.....	48
Filters, QoS, and Security.....	49
Filter Scaling.....	49
OAM and Diagnostics.....	53
Extreme Integrated Application Hosting Scaling.....	54
Fabric Scaling.....	55
Maximum Number of SPB Multicast Data I-SIDs	58
Multi-area SPB Maximums.....	58
Number of I-SIDs Supported for the Number of Configured IS-IS Interfaces and Adjacencies.....	59
Interoperability Considerations for IS-IS External Metric.....	59
Recommendations.....	60
VRF Scaling.....	60
Important Notices.....	62
ExtremeCloud IQ Support.....	62
Compatibility with ExtremeCloud IQ Site Engine.....	62
Feature-Based Licensing	62
Memory Usage.....	63
Known Issues and Restrictions.....	64
Known Issues for this Release.....	64
Restrictions and Expected Behaviors.....	84
General Restrictions and Expected Behaviors.....	85
Redirect Next-hop Filter Restrictions.....	92
Filter Restrictions.....	92
Resolved Issues this Release.....	94

Related Information.....	95
MIB Changes.....	95
Deprecated MIBs.....	95
Modified MIBs.....	95
New MIBs.....	98
Obsolete MIBs.....	102



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 supported VSP Operating System Software (VOSS) platforms.

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: 9.0.2 Release Notes revision changes

Revision	Change
AA	Initial revision for new release
AB	Added Downgrade ExtremeCloud IQ Managed Switches to 9.0.0.0 on page 21
AC	<ul style="list-style-type: none">• Updated port-based VLANs note in Layer 2 on page 37• Added Document Revision Changes section (this table)



New in this Release

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

[Other Changes](#) on page 15

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

The following platforms support VOSS 9.0.2:

- ExtremeSwitching VSP 4900 Series
- ExtremeSwitching VSP 7400 Series



Important

VSP 4450 Series, VSP 7200 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series are not supported in this release. For information about these platforms, see the VOSS 8.10.x documentation.

For MIB-related changes, see [MIB Changes](#) on page 95.



Important

VOSS 8.2 introduced changes to Segmented Management Instance that required migration of legacy management interfaces. Before you upgrade to VOSS 8.2 or later from an earlier release, you must consider your management interface configuration and migration scenario requirements. Back up and save your configuration files off the switch before upgrading to this release.

New Software Features or Enhancements

The following sections describe what is new in this release:

Channelization and MLT Support with 100G-LR4-QSFP10KM

This release adds support for channelization and MLT with 100G-LR4-QSFP10KM at both ends of the MLT link.

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

DHCP Server Enhancement

You can now configure DHCP vendor-specific options (Option 43) and apply configuration to clients that match a specific vendor class (Option 60).

For more information, see [VOSS CLI Commands Reference](#) and [VOSS User Guide](#).

General Enhancements

This release introduces the following enhancements:

- Additional resource information—The **show fulltech** command includes information for the **show ip arp spbm-tunnel-as-mac [vrf WORD<1-16>] [vrfids WORD<0-512>]** command.
- Auto-sense for Fabric Attach (FA) enhancement—You can now configure the management I-SID and C-VID values for FA proxy management and FA proxy ring management independently. In the previous release, the configured management I-SID and C-VID values were used for both FA proxy management and FA proxy ring management.
- EAP enhancement—You can use the **clear eapol eap** command to clear EAP client sessions on the switch.
- IS-IS Hello Padding enhancement—IS-IS logical-interfaces with a configured IS-IS MTU no longer automatically perform IS-IS Hello Padding, regardless of the global IS-IS Hello Padding configuration.
- MACsec enhancement—The connectivity association key name increases from 16 to 32 characters for MACsec capable platforms.
- Reset the switch—The **show software** command now displays the number of seconds remaining until the switch resets in addition to the number of days, hours, and minutes.
- SSH enhancement—You can now use either the IS-IS host name or IS-IS system ID with the SSH client on a switch to establish a connection to another switch in the SPBM cloud. The Destination Device Identifier uses the format: "fe80::xxxx:xxxx:xxxx%linuxshtobmac". The last 48 bits of the link-local address represent the IS-IS system ID, and "linuxshtobmac" is the native Linux stack interface name used by this enhancement.



Note

You cannot establish an SSH connection to the boundary node from an IS-IS remote area.

- Upgrade enhancement—If you attempt to upgrade the software on a switch with inadequate space, the switch can remove log files, if that creates sufficient space to proceed with the upgrade.
- ZTP+ enhancement—The ZTP+ agent can now configure I-SID names.

For more information, see [VOSS CLI Commands Reference](#) and [VOSS User Guide](#).

Multicast Lite RADIUS VSA Enhancement

Prior to this release, the Extreme-Dynamic-Client-Assignments Vendor Specific Attribute (VSA) failed if the VLAN and I-SID were statically configured. With this enhancement, configuration is applied even if the VLAN and I-SID are statically configured.

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

Power Reporting Enhancement

show sys-info power and **show sys power power-supply** commands now include current, voltage, and power input and output details for the power supply unit installed in your switch.

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

RADIUS Configuration using Fully Qualified Domain Name (FQDN)

In previous releases, you could only specify an IPv4 or IPv6 address when you configured a RADIUS server host or an IPv4 address for a RADIUS dynamic client. Now, you can use an FQDN, such as `host.example.com`, instead of an IP address.

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

Other Changes

Deprecated Command

The **show sys topology-ip** command is deprecated in this release, and related content is removed.

Filenames for this Release



Important

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

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

When extracting the software image file, the extraction process appends the software version portion of the extracted filenames to include the final full software version. (For example, extracting **VOSS4900.8.10.0.0.tgz** results in a software file named **VOSS4900.8.10.0.0.GA**.) Ensure that you specify the final full software version when using CLI commands that include the software version, such as activating or removing the software.

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

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

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

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

Table 5: VSP 4900 Series Software Filenames and Sizes

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
Logs reference	VOSS4900.9.0.2.0_edoc.tar	64,583,680 bytes
MD5 Checksum files	VOSS4900.9.0.2.0.md5	611 bytes
MIB - supported object names	VOSS4900.9.0.2.0_mib_sup.txt	1,548,318 bytes
MIB - objects in the OID compile order	VOSS4900.9.0.2.0_mib.txt	8,284,379 bytes
MIB - zip file of all MIBs	VOSS4900.9.0.2.0_mib.zip	1,232,811 bytes
Open source software - Master copyright file	VOSS4900.9.0.2.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	VOSS4900.9.0.2.0.sha512	1,722 bytes
Software image	VOSS4900.9.0.2.0.tgz	333,389,770 bytes
EDM Help files	VOSSv9.0.2_HELP_EDM_gzip.zip	5,234,547 bytes

Table 6: VSP 7400 Series Software Filenames and Sizes

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
Logs reference	VOSS7400.9.0.2.0_edoc.tar	64,583,680 bytes
MD5 Checksum files	VOSS7400.9.0.2.0.md5	611 bytes
MIB - supported object names	VOSS7400.9.0.2.0_mib_sup.txt	1,549,777 bytes
MIB - objects in the OID compile order	VOSS7400.9.0.2.0_mib.txt	8,284,379 bytes
MIB - zip file of all MIBs	VOSS7400.9.0.2.0_mib.zip	1,232,811 bytes
Open source software - Master copyright file	VOSS7400.9.0.2.0_oss-notice.html	2,889,456 bytes
SHA512 Checksum files	VOSS7400.9.0.2.0.sha512	1,722 bytes
Software image	VOSS7400.9.0.2.0.tgz	332,911,536 bytes
EDM Help files	VOSSv9.0.2_HELP_EDM_gzip.zip	5,234,547 bytes



Upgrade and Downgrade Considerations

[Validated Upgrade Paths](#) on page 17

[Switches That Will Not Use Zero Touch Deployment](#) on page 18

[Switches That Will Use Zero Touch Deployment](#) on page 18

[Compatible Fabric IPsec Gateway Versions](#) on page 20

[Downgrade Considerations](#) on page 20

[Migration to Segmented Management Instance](#) on page 22

[Post Upgrade Configuration for Zero Touch Fabric Configuration and Nickname Assignment](#) on page 25

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

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

- Fabric:
 - Pre-upgrade instructions for IS-IS metric type
- Considerations for VLANs or MLTs where the VLAN or MLT name uses all numbers.
- Considerations for digital certificates configured prior to VOSS 8.1.
- Considerations for Fast PoE and Perpetual PoE features configured prior to VOSS 8.1.5.

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

- [Switches That Will Not Use Zero Touch Deployment](#) on page 18
- [Switches That Will Use Zero Touch Deployment](#) on page 18

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 7: Validated upgrade paths

Product	VOSS 8.5.x to VOSS 9.0.x	VOSS 8.8.x to VOSS 9.0.x	VOSS 8.9.x to VOSS 9.0.x	VOSS 8.10.x to VOSS 9.0.x
VSP 4900 Series	Y	Y	Y	Y
VSP 7400 Series	Y	Y	Y	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. For switches prior to VOSS 8.2, migrate the Management IP address. For more information, see [Migration to Segmented Management Instance](#) on page 22 and [VOSS User Guide](#).
2. Upgrade to this release from one of the previously described releases, see [Validated Upgrade Paths](#) on page 17.
3. 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 17.
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.
- 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.
- 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 [VOSS User Guide](#).

Compatible Fabric IPsec Gateway Versions

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

For more information about image files in this release, see [Filenames for this Release](#) on page 15. For virtual service upgrade instructions, see [VOSS 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 21.

ExtremeCloud IQ Agent

For devices running VOSS 8.3, 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.

For information about how to reinstall ExtremeCloud IQ Agent firmware, beginning with VOSS 8.4.2, see [VOSS 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
```

Migration to Segmented Management Instance



Important

VOSS 8.2 introduced changes to Segmented Management Instance that required migration of legacy management interfaces. Before you upgrade to VOSS 8.2 or later from an earlier release, you must consider your management interface configuration and migration scenario requirements. Backup and save your configuration files off the switch before upgrading to this release.

If the switch already runs VOSS 8.2 or later, you can ignore this section.

Management interface access to the switch can be lost if you do not perform the applicable migration scenarios before upgrading to this release. Loss of management access after an upgrade can result in an automatic roll-back to the previous software version.

You must perform a manual software commit after upgrading from VOSS Release 8.1.5.0 or earlier to VOSS 8.2 or later. Management interface access is required to input the **software commit** CLI command within 10 minutes after the upgrade. If the time expires the system initiates an automatic roll-back to the previous release.

You must ensure the switch runs VOSS 8.1.x before you upgrade to VOSS 8.2 or later to support the **migrate-to-mgmt** functionality.



Note

If the network environment must migrate static IPv6 routes, the switches must run VOSS Release 8.1.2.0 or later before you upgrade to VOSS 8.2 or later.

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 17.

You must consider the following legacy management interface migration scenarios before you upgrade to VOSS 8.2 or later:

Table 8: Management Interface Migration Scenarios

Mgmt Interface	Mgmt Scenario	Migration Description
DvR leaf	Automatic migration during upgrade.	DvR leaf settings migrate automatically during the software upgrade process. The DvR inband-mgmt-ip CLIP automatically becomes the new Segmented Management Instance CLIP. Note: Leaf nodes only support the management CLIP as part of the Global Routing Table (GRT).
OOB	Automatic migration during upgrade.	Out-of-Band management settings migrate automatically during the software upgrade process.

Table 8: Management Interface Migration Scenarios (continued)

Mgmt Interface	Mgmt Scenario	Migration Description
CLIP	Specify a Circuitless IP (CLIP) interface for migration to management interface before you upgrade.	<p>You can use this interface type for CLIP management network routing in a Fabric network or Layer 3 routing network.</p> <p>Use the migrate-to-mgmt command in the Loopback Interface Configuration mode of the CLI to specify the CLIP interface for management before starting the software upgrade process.</p> <p>You can designate the IP Shortcut CLIP to migrate to the Management Instance CLIP. After the upgrade, the IS-IS source IP address moves to the Management Instance CLIP. You should configure a new GRT CLIP using a different IP address and assign that as the new IS-IS source IP. Save the configuration before upgrading.</p> <p>Important: Ensure that the management CLIP IP address does not fall into the range of a configured VLAN IP address range as this is not allowed.</p>
VLAN	Specify a VLAN interface for migration to management interface before you upgrade.	<p>You can use this interface type for management of Layer 2 switches or for Zero-Touch onboarding of newly deployed devices. Use the CLIP Management Instance for routed management.</p> <p>Use the migrate-to-mgmt command in the VLAN Interface Configuration mode of the CLI to specify the VLAN interface for management before starting the software upgrade process.</p> <p>Important: Choose a VLAN that does not have an IP interface on it. The upgrade process removes the</p>

Table 8: Management Interface Migration Scenarios (continued)

Mgmt Interface	Mgmt Scenario	Migration Description
		<p>IP configuration and network connectivity can be impacted.</p> <p>Save the configuration before upgrading.</p> <p>The VLAN Management Instance does not route to or from the GRT. Bridged management traffic must ingress on the VLAN or I-SID.</p>

For more information about Segmented Management Instance migration, see [VOSS User Guide](#).

Segmented Management Instance Migration and DvR

Starting with VOSS Release 8.2, VSP devices can be managed by a CLIP/Loopback IP address that is assigned to a virtual router and forwarder (VRF) that is not in the Global Routing Table (GRT). When you convert a VSP switch from a regular backbone edge bridge (BEB) to a DvR leaf device by setting the DvR leaf boot flag, you must assign the management CLIP to the GRT. If you assign the management CLIP to a VRF, the device will not be reachable after the migration because the management CLIP cannot be migrated.

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 VOSS 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 VOSS 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 [VOSS 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 17.

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 VOSS 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 15999999.
 - If the DHCP server is on the same subnet as the onboarding I-SID, configure the port facing the DHCP server as private-vlan promiscuous, using Private VLAN 4048, if the new DHCP snooping port feature does not have the promiscuous port configured automatically. This VLAN and the Auto-sense onboarding I-SID are created automatically on a newly deployed device.
- In this release, ports send Fabric Connect LLDP TLVs regardless of the Auto-sense configuration, which means these devices can establish adjacencies with other devices that run a Zero Touch Fabric release, and use either Auto-sense or static NNI configuration.

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

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

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

Zero Touch Fabric Configuration Switch



Important

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

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

1. Upgrade to a Zero Touch Fabric release, if the device is not a new deployment already running a Zero Touch Fabric release.
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 [VOSS 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 VOSS 8.10.

5. The switch joins the Fabric.
6. For devices running releases earlier than VOSS 8.10, the nickname server dynamically assigns an SPBM nickname. For devices running releases VOSS 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

- [VSP 4450 Series Hardware](#) on page 29
- [VSP 4900 Series Hardware](#) on page 29
- [VSP 7200 Series Hardware](#) on page 32
- [VSP 7400 Series Hardware](#) on page 33
- [VSP 8000 Series Hardware](#) on page 34
- [XA1400 Series Hardware](#) on page 34
- [Transceivers](#) on page 34
- [Power Supply Compatibility](#) on page 35

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

VSP 4450 Series Hardware



Important

For information related to VSP 4450 Series, VSP 7200 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series, see the VOSS 8.10.x documentation.

VSP 4900 Series Hardware

Table 9: Switch models

Model	Initial release	Supported new VOSS feature release				
		8.9	8.10	8.10.1	9.0	9.0.2
VSP4900-48P	8.1	Y	Y	Y	Y	Y
VSP4900-12MXU-12XE	8.1.5	Y	Y	Y	Y	Y

Table 9: Switch models (continued)

Model	Initial release	Supported new VOSS feature release				
		8.9	8.10	8.10.1	9.0	9.0.2
VSP4900-24S	8.1.5	Y	Y	Y	Y	Y
VSP4900-24XE	8.1.5	Y	Y	Y	Y	Y

**Note**

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

Table 10: Versatile Interface Modules (VIM)

Model	Initial release	Supported new VOSS feature release				
		8.9	8.10	8.10.1	9.0	9.0.2
VIM5-4X	8.1	Y	Y	Y	Y	Y
VIM5-4XE	8.1	Y	Y	Y	Y	Y
VIM5-2Y	8.1	Y	Y	Y	Y	Y
VIM5-4YE	8.1	Y	Y	Y	Y	Y
VIM5-2Q	8.1	Y	Y	Y	Y	Y
VIM5-4Y	8.1.5	Y	Y	Y	Y	Y

VSP 4900 Series Operational Notes

VSP4900-24S fixed ports operate at 1 Gbps. If you connect a 10 Gbps DAC/SFP+ to a VSP4900-24S 1 Gbps fixed port, the system displays the following error message:

10Gb optical module inserted in 1Gb only port nn. Not supported.

Although the link successfully comes up, the operational speed shows as 10 Gbps instead of 1 Gbps. This scenario occurs when a 10 Gbps DAC/SFP+ is used to make any of the following connections from a VSP4900-24S 1 Gbps fixed port:

- a VSP4900-24S to VSP4900-24S loopback connection
- a VSP4900-24S connected to another VSP4900-24S
- a VSP4900-24S connected to a VSP 4450GSX

Versatile Interface Module Operational Notes

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

Table 11: VSP 4900 Series VIM Matrix

	VIM5-4X	VIM5-4XE	VIM5-2Y	VIM5-4YE	VIM5-4Y	VIM5-2Q
Number of supported ports for VSP4900-48P and VSP4900-24S	4	4	2	2	2	1
Number of supported ports for VSP4900-24XE and VSP4900-12MXU-12XE	4	4	2	4	4	2
Port speeds	1 Gbps 10 Gbps	1 Gbps 10 Gbps	10 Gbps or 25 Gbps All ports must operate at either 10 Gbps or 25 Gbps (default)	10 Gbps or 25 Gbps All ports must operate at either 10 Gbps or 25 Gbps (default)	10 Gbps or 25 Gbps All ports must operate at either 10 Gbps or 25 Gbps (default)	40 Gbps 10 Gbps (with channeliz ation
PHY present	No	Yes	Yes	Yes	Yes	No
Copper transceiver support (1 Gbps/10 Gbps)	10GBASE-T only	Both	10GBASE-T only	10GBASE-T only	10GBASE-T only	Not applicabl e
MACsec	Not supported	128/256 bit	Not supporte d	128/256 bit	Not supporte d	Not supporte d
Forward Error Correction (FEC)	Not supported	Not supporte d	Not supporte d	Default is Auto-FEC - FEC Auto, CL108, CL91, CL74 and No FEC supporte d	Not supporte d	Not supporte d
1 Gbps Auto-Negotiation	Disabled	Enabled	Not applicabl e	Not applicabl e	Not applicabl e	Not applicabl e
10 Gbps Auto-Negotiation	Disabled	Disabled	Disabled	Disabled	Disabled	Not applicabl e

Table 11: VSP 4900 Series VIM Matrix (continued)

	VIM5-4X	VIM5-4XE	VIM5-2Y	VIM5-4YE	VIM5-4Y	VIM5-2Q
25 Gbps Auto-Negotiation	Not applicable	Not applicable	Disabled	Enabled for DACs Disabled for AOCs, optical transceivers	Disabled	Not applicable
<p>Note: Auto-Negotiation values are automatically set based on the type of transceiver detected.</p>						

VIM5-2Y and VIM5-4Y Operational Notes



Note

VIM5-2Y and VIM5-4Y are in end-of-sale status.

The IEEE 802.3by requirement for 25 G is that any transceiver or DAC 3 meters or longer, requires the use of forward error correction (FEC). Because the VIM5-2Y and VIM5-4Y do not support FEC, note the following considerations for proper operation with these VIMs:

- Supported 25 G optics:
 - PN: 10502 - 25GBASE-SR (FEC-Lite): up to 30 m for OM3, up to 40 m for OM4
- Supported 25 G DACs:
 - 10520 25G SFP28 Cable (1 m)
- You must disable Auto-Negotiation and FEC on any VSP 7400 Series device that is connected to either of these VIMs.

You might experience CRC or link flap errors by using an unsupported 25 G transceiver.

VSP 7200 Series Hardware



Important

For information related to VSP 4450 Series, VSP 7200 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series, see the VOSS 8.10.x documentation.

VSP 7400 Series Hardware

Part number	Model Number	Initial release	Supported new VOSS feature release				
			8.9	8.10	8.10.1	9.0	9.0.2
VSP7400-32C (no power supplies or fans) VSP7400-32C-AC-F (front-to-back airflow) VSP7400-32C-AC-R (back-to-front airflow)	VSP 7432CQ	8.0	Y	Y	Y	Y	Y
VSP7400-48Y-8C (no power supplies or fans) VSP7400-48Y-8C-AC-F (front-to-back airflow) VSP7400-48Y-8C-AC-R (back-to-front airflow)	VSP 7400-48Y	8.0.5	Y	Y	Y	Y	Y

VSP 7400 Series Operational Notes

The VSP 7400 Series has a PHYless design. The benefits of a PHYless design are lower power consumption and lower latency. However, due to the PHYless design, some transceivers that require electronic dispersion compensation (EDC) for proper operation are not supported. For a list of supported transceivers, see [Extreme Optics](#) website.

The following list provides operational notes for VSP 7432CQ.

- Ports 31 and 32 (low) or ports 29, 30, 31, and 32 (high) are reserved for internal use when certain features, including Fabric Connect, are used. For a full list of the features, refer to [VOSS User Guide](#).
- The QSFP28 ports support the use of QSFP28 and QSFP+ transceivers:
 - The software detects the transceiver type and sets the port speed as either 100 Gbps for QSFP28 or 40 Gbps for QSFP+.
- Channelization:
 - Channelization is not supported on port 28.
 - Supports 4x10 Gbps when channelization is enabled and QSFP+ transceiver is detected.
 - Supports 4x25 Gbps when channelization is enabled and QSFP28 transceiver is detected.

The following list provides operational notes for VSP 7400-48Y.

- Ports 55 and 56 (low) or ports 53, 54, 55, and 56 (high) are reserved for internal use when certain features, including Fabric Connect, are used. For a full list of the features, refer to [VOSS User Guide](#).
- The QSFP28 ports support the use of QSFP28 and QSFP+ transceivers:
 - The software detects the transceiver type and sets the port speed as either 100 Gbps for QSFP28 or 40 Gbps for QSFP+.
- The SFP28 ports support the use of SFP28, SFP, and SFP+ transceivers.
 - The software detects the transceiver type and sets the port speed as either 25 Gbps for SFP28, 1 Gbps for SFP, or 10 Gbps for SFP+.
 - Auto-Negotiation is not supported when a 25 Gbps port operates at 1 Gbps. The following log message displays on the switch: `Auto-Negotiation enabled but not applied to port 1/1 since 1G transceiver is present..`
- Channelization is not supported. As a result, you cannot use the following optical components:
 - 40 Gbps or 100 Gbps breakout cables
 - QSFP28 to SFP28 Adapter (PN: 10506)

VSP 8000 Series Hardware



Important

For information related to VSP 4450 Series, VSP 7200 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series, see the VOSS 8.10.x documentation.

XA1400 Series Hardware



Important

For information related to VSP 4450 Series, VSP 7200 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series, see the VOSS 8.10.x documentation.

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 [VOSS 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:

- [VSP 4900 Series Switches: Hardware Installation Guide](#)
- [VSP 7400 Series Switches: Hardware Installation Guide](#)



Scaling

[Layer 2](#) on page 37

[IP Unicast](#) on page 40

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

[IP Multicast](#) on page 45

[Distributed Virtual Routing \(DvR\)](#) on page 47

[VXLAN Gateway](#) on page 48

[Filters, QoS, and Security](#) on page 49

[OAM and Diagnostics](#) on page 53

[Extreme Integrated Application Hosting Scaling](#) on page 54

[Fabric Scaling](#) on page 55

[VRF Scaling](#) on page 60

This section documents scaling capabilities of the VOSS 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 [VOSS User Guide](#).

Layer 2

Table 12: Layer 2 Maximums

Attribute	Product	Maximum number supported
MAC table size (without SPBM)	VSP 4900 Series	80,000
	VSP 7400 Series	160,000
MAC table size (with SPBM)	VSP 4900 Series	40,000
	VSP 7400 Series	80,000
Endpoint Tracking MAC addresses per switch	VSP 4900 Series	8,000
	VSP 7400 Series	8,000
Directed Broadcast interfaces	VSP 4900 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 39.
	VSP 7400 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 39.
Port-based VLANs Note: When you use Flex-UNI functionality, you can use the range from 1 to 4094 for port VLAN IDs.	VSP 4900 Series	4,059
	VSP 7400 Series	4,059
Private VLANs	VSP 4900 Series	200
	VSP 7400 Series	200
Protocol-based VLANs (IPv6 only)	VSP 4900 Series	1
	VSP 7400 Series	1
RSTP instances	VSP 4900 Series	1
	VSP 7400 Series	1
MSTP instances	VSP 4900 Series	12
	VSP 7400 Series	64

Table 12: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
LACP aggregators	VSP 4900 Series	VSP4900-48P: 52 (48 fixed ports + 4 VIM ports) VSP4900-24S, VSP4900-24XE, VSP4900-12MXU-12XE: 28 (24 fixed + 4 VIM ports)
	VSP 7400 Series	VSP 7432CQ: 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y: 56 configured in Full Port mode
Ports per LACP aggregator	VSP 4900 Series	8 active
	VSP 7400 Series	8 active
MLT groups	VSP 4900 Series	VSP4900-48P: 52 (48 fixed ports + 4 VIM ports) VSP4900-24S, VSP4900-24XE, VSP4900-12MXU-12XE: 28 (24 fixed + 4 VIM ports)
	VSP 7400 Series	VSP 7432CQ: 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y: 56 configured in Full Port mode
Ports per MLT group	VSP 4900 Series	8
	VSP 7400 Series	8
Link State Tracking (LST) groups	VSP 4900 Series	48
	VSP 7400 Series	48
Interfaces per LST group	VSP 4900 Series	8 upstream 128 downstream
	VSP 7400 Series	8 upstream 128 downstream
SLPP VLANs	VSP 4900 Series	128
	VSP 7400 Series	500

Table 12: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
VLACP interfaces	VSP 4900 Series	VSP4900-48P: 52 (48 fixed ports + 4 VIM ports) VSP4900-24S, VSP4900-24XE, VSP4900-12MXU-12XE: 28 (24 fixed + 4 VIM ports) VIM5-2Q on VSP4900-12MXU-12XE and VSP4900-24XE with channelization enabled: 32
	VSP 7400 Series	VSP 7432CQ : 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y: 56 configured in Full Port mode
Microsoft NLB cluster IP interfaces	VSP 4900 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 39.
	VSP 7400 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 39.

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 13: IP Unicast Maximums

Attribute	Product	Maximum number supported
IP interfaces (IPv4 or IPv6 or IPv4+IPv6)	VSP 4900 Series	500 See IP Interface Maximums for VSP 4900 Series on page 43.
	VSP 7400 Series	1,000 See IP Interface Maximums for VSP 7400 Series on page 43.
VRRP interfaces (IPv4 or IPv6)	VSP 4900 Series	252 See IP Interface Maximums for VSP 4900 Series on page 43.
	VSP 7400 Series	500 per switch 256 per VRF See IP Interface Maximums for VSP 7400 Series on page 43.
Routed Split Multi-Link Trunking (RSMLT) interfaces (IPv4 or IPv6 or IPv4+IPv6)	VSP 4900 Series	251 See IP Interface Maximums for VSP 4900 Series on page 43.
	VSP 7400 Series	499 See IP Interface Maximums for VSP 7400 Series on page 43.
VRRP interfaces with fast timers (200ms) - IPv4/IPv6	VSP 4900 Series	24
	VSP 7400 Series	24
ECMP groups/paths per group	VSP 4900 Series	2,048/8
	VSP 7400 Series	2,048/8
OSPF v2/v3 interfaces	VSP 4900 Series	500
	VSP 7400 Series	500
OSPF v2/v3 neighbors (adjacencies)	VSP 4900 Series	500
	VSP 7400 Series	500

Table 13: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
OSPF areas	VSP 4900 Series	12 for each VRF 80 for the switch
	VSP 7400 Series	12 for each VRF 80 for the switch
IPv4 ARP table	VSP 4900 Series	32,000 in non-SPB deployments 16,000 in SPB deployments
	VSP 7400 Series	56,000 non-SPB deployments 40,000 SPB deployments
IPv4 CLIP interfaces	VSP 4900 Series	64
	VSP 7400 Series	64
IPv4 RIP interfaces	VSP 4900 Series	200
	VSP 7400 Series	200
IPv4 BGP peers	VSP 4900 Series	256
	VSP 7400 Series	256
IPv4 VRFs with iBGP	VSP 4900 Series	16
	VSP 7400 Series	16
IPv4/IPv6 VRF instances For additional information, see VRF Scaling on page 60.	VSP 4900 Series	256 including mgmt VRF and GRT See IP Interface Maximums for VSP 4900 Series on page 43.
	VSP 7400 Series	256 including mgmt VRF and GRT See IP Interface Maximums for VSP 7400 Series on page 43.
IPv4 static ARP entries	VSP 4900 Series	2,000 for each VRF 10,000 for the switch
	VSP 7400 Series	2,000 for each VRF 10,000 for the switch
IPv4 static routes	VSP 4900 Series	1,000 for each VRF 5,000 for the switch
	VSP 7400 Series	1,000 for each VRF 5,000 for the switch

Table 13: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
IPv4 route policies	VSP 4900 Series	500 for each VRF 5,000 for the switch
	VSP 7400 Series	500 for each VRF 5,000 for the switch
IPv4 UDP forwarding entries	VSP 4900 Series	512
	VSP 7400 Series	1,024
IPv4 DHCP Relay forwarding entries	VSP 4900 Series	2,048
	VSP 7400 Series	2,048
IPv6 DHCP Snoop entries in Source Binding Table	VSP 4900 Series	1,024
	VSP 7400 Series	1,024
IPv6 Neighbor table	VSP 4900 Series	8,000
	VSP 7400 Series	32,000
IPv6 static entries in Source Binding Table	VSP 4900 Series	256
	VSP 7400 Series	256
IPv6 static neighbor records	VSP 4900 Series	128 per VRF 512 per system
	VSP 7400 Series	128 per VRF 512 per system
IPv6 CLIP interfaces	VSP 4900 Series	64
	VSP 7400 Series	64
IPv6 static routes	VSP 4900 Series	1,000
	VSP 7400 Series	1,000
IPv6 6in4 configured tunnels	VSP 4900 Series	64
	VSP 7400 Series	64
IPv6 DHCP Relay forwarding	VSP 4900 Series	512 per switch 10 per VRF
	VSP 7400 Series	512
IPv6 BGP peers	VSP 4900 Series	256 Up to 8,000 IPv6 prefixes for BGPv6 peering
	VSP 7400 Series	256
IPv6 VRFs with iBGP	VSP 4900 Series	16
	VSP 7400 Series	16
BFD VRF instances	VSP 4900 Series	16
	VSP 7400 Series	16

Table 13: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
BFD sessions per switch (IPv4/IPv6) with default values	VSP 4900 Series	16
	VSP 7400 Series	16
BFD sessions per switch (IPv4) with 750ms timers for BGP and static routes only	VSP 4900 Series	16
	VSP 7400 Series	50
BFD sessions with Fabric Extend tunnels (IPv4)	VSP 4900 Series	16
	VSP 7400 Series	16

IP Interface Maximums for VSP 4900 Series

The maximum number of IP interfaces for VSP 4900 Series is based on the following formulas:

- If you disable the VRF scaling boot configuration flag:
 - = 500 – (# of VRRP IPv4 interfaces) - (# of VRRP IPv6 interfaces) – (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3x(# of VRFs)
- If you enable the VRF scaling boot configuration flag:
 - = 500 – (# of VRRP IPv4 interfaces) – (# of VRRP IPv6 interfaces) - (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3

IP Interface Maximums for VSP 7400 Series

The maximum number of IP interfaces for VSP 7400 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 with unique mac offset + (# 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 with unique mac offset + 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 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

Layer 3 Route Table Size

Table 14: Layer 3 Route Table Size Maximums

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

Route Scaling

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

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

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

Table 15: VSP 4900 Series

URPF mode	IPv6 mode	IPv4	IPv6	
			Prefix less than 64	Prefix greater than 64
No	No	15,488	7,744	n/a
No	Yes	7,488	3,744	2,000
Yes	No	7,488	3,744	n/a

Table 15: VSP 4900 Series (continued)

URPF mode	IPv6 mode	IPv4	IPv6	
			Prefix less than 64	Prefix greater than 64
Yes	Yes	3,488	1,744	2,000

Note:
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 OR 'z' ipv6 >64 routes (not a combination of all).

Table 16: VSP 7400 Series

URPF mode	IPv6 mode	IPv4	IPv6	
			Prefix less than 64	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 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 OR 'z' ipv6 >64 routes (not a combination of all).

IP Multicast

Table 17: IP Multicast Maximums

Attribute	Product	Maximum number supported
IGMP/MLD interfaces (IPv4/IPv6)	VSP 4900 Series	4,059
	VSP 7400 Series	4,059
PIM interfaces (IPv4/IPv6)	VSP 4900 Series	128 Active
	VSP 7400 Series	128 Active
PIM Neighbors (IPv4/IPv6) (GRT Only)	VSP 4900 Series	128
	VSP 7400 Series	128
PIM-SSM static channels (IPv4/IPv6)	VSP 4900 Series	4,000
	VSP 7400 Series	4,000

Table 17: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
Multicast receivers/IGMP joins (IPv4/IPv6) (per switch)	VSP 4900 Series	6,000
	VSP 7400 Series	6,000
Total multicast routes (S,G,V) (IPv4/IPv6) (per switch)	VSP 4900 Series	6,000
	VSP 7400 Series	6,000
Total multicast routes (S,G,V) (IPv4) on an SPB-PIM Gateway configured switch	VSP 4900 Series	3,000
	VSP 7400 Series	3,000
Static multicast routes (S,G,V) (IPv4/IPv6)	VSP 4900 Series	4,000
	VSP 7400 Series	4,000
Multicast enabled Layer 2 VSN (IPv4)	VSP 4900 Series	2,000
	VSP 7400 Series	2,000
Multicast enabled Layer 3 VSN (IPv4)	VSP 4900 Series	256 including mgmt VRF and GRT
	VSP 7400 Series	256 including mgmt VRF and GRT
SPB-PIM Gateway controller S,Gs (source announcements) with MSDP (IPv4)	VSP 4900 Series	6,000
	VSP 7400 Series	6,000
SPB-PIM Gateway controllers per SPB fabric (IPv4)	VSP 4900 Series	5
	VSP 7400 Series	5
SPB-PIM Gateway nodes per SPB fabric (IPv4)	VSP 4900 Series	64
	VSP 7400 Series	64
SPB-PIM Gateway interfaces per BEB (IPv4)	VSP 4900 Series	64
	VSP 7400 Series	64
PIM neighbors per SPB-PIM Gateway node (IPv4)	VSP 4900 Series	64
	VSP 7400 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 40.

Table 18: 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 the VSP 4450 Series controls the scaling of the DvR domain it is in. For VSP 4450 Series scaling information, see VOSS Release Notes for VOSS Release 8.10. 		
DvR Virtual IP interfaces	VSP 4900 Series	499 with vIST 500 without vIST
	VSP 7400 Series	999 with vIST as interior node 1,000 without vIST as interior node 500 on boundary node
DvR domains per SPB fabric	VSP 4900 Series	16
	VSP 7400 Series	16
Controller nodes per DvR domain with default route inject flag enabled Total number of Controllers per domain cannot exceed 8.	VSP 4900 Series	8
	VSP 7400 Series	8
Leaf nodes per DvR domain	VSP 4900 Series	250
	VSP 7400 Series	250
DvR enabled Layer 2 VSNS	VSP 4900 Series	501 with vIST 502 without vIST
	VSP 7400 Series	999 with vIST 1,000 without vIST

Table 18: 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.	VSP 4900 Series	32,000
	VSP 7400 Series	40,000

VXLAN Gateway

Table 19: VXLAN Gateway Maximums

Attribute	Product	Maximum number supported
MAC addresses in base interworking mode	VSP 4900 Series	n/a
	VSP 7400 Series	80,000
MAC addresses in full interworking mode	VSP 4900 Series	n/a
	VSP 7400 Series	50,000
VNI IDs per node	VSP 4900 Series	n/a
	VSP 7400 Series	2,000
VTEP destinations per node or VTEP	VSP 4900 Series	n/a
	VSP 7400 Series	500

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

Table 20: OVSDB protocol support for VXLAN Gateway Maximums

Attribute	Product	Maximum number supported
Maximum controllers to which a single VTEP switch can connect	VSP 4900 Series	n/a
	VSP 7400 Series	3

Filters, QoS, and Security

Table 21: Filters, QoS, and Security Maximums

Attribute	Product	Maximum number supported
For more information, see Filter Scaling on page 49.		
Total IPv4 Ingress rules/ACEs (Port/VLAN/InVSN based, Security/QoS filters)	VSP 4900 Series	1,536
	VSP 7400 Series	767 Primary Bank 767 Secondary Bank
Total IPv4 Egress rules/ACEs (Port based, Security filters)	VSP 4900 Series	248
	VSP 7400 Series	783 271 if you enable <code>boot config flags ipv6-egress-filter</code>
Total IPv6 Ingress rules/ACEs (Port/VLAN/InVSN based, Security filters)	VSP 4900 Series	1024
	VSP 7400 Series	767
Total IPv6 egress rules/ACEs (Port based, Security filters)	VSP 4900 Series	256
	VSP 7400 Series	511
EAP (clients per port) Note: The total of EAP clients plus NEAP clients per port or per switch cannot exceed 8,192.	VSP 4900 Series	32
	VSP 7400 Series	32
NEAP Note: The total of EAP clients plus NEAP clients per port or per switch cannot exceed 8,192.	VSP 4900 Series	8,192 for NEAP
	VSP 7400 Series	8,192 for NEAP

Filter Scaling

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

VSP 4900 Series

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
 - 512 ACLs with 1 security ACE each OR
 - 256 ACLs with 1 QoS ACE each OR
 - a combination based on the following rule:
 - ((num ACLs + num security ACEs) <= 1024) && ((num ACLs + num QoS ACEs) <= 512)

This maximum implies a VLAN member count of 1 for inVlan ACLs

- 512 IPv6 ingress ACLs (inPort):
 - 512 ACLs with 1 security ACE each OR
 - a combination based on the following rule:
 - $(\text{num ACLs} + \text{num security ACEs}) \leq 512$
- 124 egress ACLs (outPort only):
 - 124 ACLs with 1 security 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.

- 1534 ingress ACEs:

Theoretical maximum of 1534 implies 1 ingress ACL with 1023 security ACEs and 511 QoS ACEs

- Ingress ACEs supported: $(1024 (\text{security}) - \# \text{ of ACLs}) + (512 (\text{QoS}) - \# \text{ of ACLs})$.

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

- 247 egress ACEs:

Theoretical maximum of 247 implies 1 egress ACL with 247 security ACEs

- Egress ACEs supported: $248 - \# \text{ of ACLs}$.

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

VSP 7400 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 in VSN, 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
- $511 - \text{num IPv6 egress ACLs}$

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 22: 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)
VSP 4900 Series	4	30	97	49
VSP 7400 Series	4	50	532	20

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 VSP 7400 Series for ten Routed Private VLANs with four private trunk members each.

```
Switch:1>show io resources filter
=====
                        FILTER TABLE
=====

ACL Filter Resource Manager stats
-----
BCM CAP Group: | ICAP_SEC | ICAP_QOS | ICAP_IPv6 | ECAP_SEC | ECAP_IPv6
Group Mode:   | Double   | Triple   | Triple     | Double    | Double
-----
Total Entries : | 767     | 767     | 767       | 782      | 512
Free Entries  : | 767     | 767     | 767       | 732      | 512
In Use       : | 0       | 0       | 0         | 50       | 0
Filter table:
-----
ACL |          | Port/Vlan | Sec | QoS | All |
ID | Flags | Members | ACE's | ACE's | ACE's | Type
-----

Filter resources used by other features:
-----
Feature | Type | Number of entries |
-----
Pvlan  | ECAP | 50                |
-----
```

OAM and Diagnostics

Table 23: OAM and Diagnostics Maximums

Attribute	Product	Maximum number supported
EDM sessions	VSP 4900 Series	5
	VSP 7400 Series	5
FTP sessions (IPv4/IPv6)	VSP 4900 Series	8 total (4 for IPv4 and 4 for IPv6)
	VSP 7400 Series	8 total (4 for IPv4 and 4 for IPv6)
SSH sessions (IPv4/IPv6)	VSP 4900 Series	8 total (any combination of IPv4 and IPv6)
	VSP 7400 Series	8 total (any combination of IPv4 and IPv6)
Telnet sessions (IPv4/IPv6)	VSP 4900 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 7400 Series	16 total (8 for IPv4 and 8 for IPv6)
TFTP sessions (IPv4/IPv6)	VSP 4900 Series	2 total (any combination of IPv4 and IPv6)
	VSP 7400 Series	2 total (any combination of IPv4 and IPv6)
Mirrored ports (source)	VSP 4900 Series	51 (52 ports per chassis, 48 fixed ports plus up to 4 ports on the VIMs)
	VSP 7400 Series	31 (up to 125 with channelization) with Advanced Feature Bandwidth Reservation configured in Full Port mode
Mirroring ports (destination)	VSP 4900 Series	4
	VSP 7400 Series	4
Fabric RSPAN Port mirror instances per switch (Ingress only)	VSP 4900 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	VSP 7400 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 23: OAM and Diagnostics Maximums (continued)

Attribute	Product	Maximum number supported
Fabric RSPAN Flow mirror instances per switch (Ingress only)	VSP 4900 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	VSP 7400 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
Fabric RSPAN Monitoring I-SIDs (network value)	VSP 4900 Series	1,000 Monitoring I-SIDs across SPB network
	VSP 7400 Series	1,000 Monitoring I-SIDs across SPB network
sFlow sampling limit	VSP 4900 Series	3,100 samples per second
	VSP 7400 Series	9,000 samples per second
IPFIX flows	VSP 4900 Series	n/a
	VSP 7400 Series	32,767
Application Telemetry host monitoring - maximum number of monitored hosts	VSP 4900 Series	382 hosts
	VSP 7400 Series	767 hosts
Note: These resources are shared with the IPv4 Filter Ingress rules/ACEs.		

Extreme Integrated Application Hosting Scaling

Table 24: Extreme Integrated Application Hosting (IAH) Maximums

Attribute	Product	Maximum number supported
Simultaneous Virtual Machines	VSP 4900 Series	Not supported
	VSP 7400 Series	6
CPU cores available to VMs	VSP 4900 Series	2
	VSP 7400 Series	6
Memory available to VMs	VSP 4900 Series	4 GB
	VSP 7400 Series	12 GB
Storage available to VMs	VSP 4900 Series	104 GB of 120 modular SSD
	VSP 7400 Series	100 GB
Total SRIOV vports available to VMs	VSP 4900 Series	16
	VSP 7400 Series	16

Table 24: Extreme Integrated Application Hosting (IAH) Maximums (continued)

Attribute	Product	Maximum number supported
Vports available to single VM	VSP 4900 Series	16
	VSP 7400 Series	16

Fabric Scaling

This section lists the fabric scaling information.

Table 25: Fabric Maximums

Attribute	Product	Maximum number supported (with and without vIST)
Number of SPB IS-IS areas	VSP 4900 Series	1
	VSP 7400 Series as Interior Node	1
	VSP 7400 Series as Boundary Node	2
Number of B-VIDs	VSP 4900 Series	2
	VSP 7400 Series	2
Maximum number of Physical and Logical (Fabric Extend) NNI interfaces/adjacencies (Home and Remote area total when operating as Boundary Node)	VSP 4900 Series	255, of which 64 can be with IPsec using Fabric IPsec Gateway
	VSP 7400 Series	255, of which 64 can be with IPsec using Fabric IPsec Gateway
SPBM enabled nodes per area (BEB + BCB)	VSP 4900 Series	800
	VSP 7400 Series as Interior Node	2,000
	VSP 7400 Series as Boundary Node	500 per area
Number of BEBs not part of vIST clusters this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI)	VSP 4900 Series	500
	VSP 7400 Series	2,000
Number of BEBs that are part of a vIST cluster this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI)	VSP 4900 Series	330
	VSP 7400 Series	1,330
I-SIDs supported (local UNI present on device)	VSP 4900 Series	See Number of I-SIDs supported
	VSP 7400 Series	See Number of I-SIDs supported

Table 25: Fabric Maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
I-SIDs supported on Boundary Nodes (no local UNI present on device)	VSP 4900 Series	n/a
	VSP 7400 Series as Boundary Node	9,600
Maximum number of Layer 2 VSNs per switch (local UNI present on device)	VSP 4900 Series	4,059
	VSP 7400 Series	4,000
Maximum number of inter-area redistributed Layer 2 VSNs (no local UNI present on Boundary Node)	VSP 4900 Series	n/a
	VSP 7400 Series as Boundary Node	9,600
Maximum number of Switched UNI Endpoints (C-VID or untagged port bindings)	VSP 4900 Series	8,000
	VSP 7400 Series	12,000
Maximum number of Transparent Port UNIs per switch	VSP 4900 Series	52
	VSP 7400 Series	VSP 7432CQ: 30 (up to 120 with channelization) configured in Full Port mode VSP 7400-48Y: 54 configured in Full Port mode
Maximum number of E-Tree PVLAN UNIs per switch	VSP 4900 Series	200
	VSP 7400 Series	200
Maximum number of Layer 3 VSNs per switch See VRF Scaling on page 60.	VSP 4900 Series	256 including mgmt VRF and GRT
	VSP 7400 Series	256 including mgmt VRF and GRT
Maximum number of SPB Layer 2 multicast Data I-SIDs	VSP 4900 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 58
	VSP 7400 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 58

Table 25: Fabric Maximums (continued)

Attribute	Product	Maximum number supported (with and without vIST)
Maximum number of SPB Layer 3 multicast Data I-SIDs	VSP 4900 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 58 Note: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached.
	VSP 7400 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 58 Note: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached.
Maximum number of FA ISID/VLAN assignments per port	VSP 4900 Series	94
	VSP 7400 Series	94
Maximum number of IP multicast S,Gs when operating as a BCB (intra-area)	VSP 4900 Series	16,000
	VSP 7400 Series	50,000
Maximum number of IP multicast S,Gs when operating as a Boundary Node (inter-area)	VSP 4900 Series	n/a
	VSP 7400 Series as Boundary Node	4,800
ISW switches in a Fabric Attach Ring		128

Maximum Number of SPB Multicast Data I-SIDs

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

Attribute		Product	Maximum number supported (with and without vIST)
Maximum number of SPB Layer 2 multicast Data I-SIDs Note: Overall limits across Layer 2 VSNs	On Ingress BEB: Dynamic and Static originated Data I-SIDs	VSP 4900 Series	4,000
		VSP 7400 Series as Boundary Node	4,000
	On Egress BEB: Static Data I-SIDs Terminated	VSP 4900 Series	6,000
		VSP 7400 Series as Boundary Node	6,000
	On Egress BEB: Dynamic data I-SIDs + originating BEB pairs terminated	VSP 4900 Series	6,000
		VSP 7400 Series as Boundary Node	6,000
Maximum number of SPB Layer 3 multicast Data I-SIDs Note: Overall limits across all Layer 3 VSNs/GRT	On Ingress BEB: Dynamic and Static originated Data I-SIDs	VSP 4900 Series	4,000
		VSP 7400 Series as Boundary Node	4,000
	On Egress BEB: Static Data I-SIDs Terminated	VSP 4900 Series	6,000
		VSP 7400 Series as Boundary Node	6,000
	On Egress BEB: Dynamic data I-SIDs + originating BEB pairs terminated	VSP 4900 Series	6,000
		VSP 7400 Series as Boundary Node	6,000

Multi-area SPB Maximums

Table 26: Multi-area SPB maximums

Scaling	VSP 7400 Series
SPBM enabled nodes per area	500
SPBM total nodes home + remote	1,000
I-SIDs supported on boundary nodes (no local UNI present on device)	9,600
Maximum number of inter-area redistributed Layer 2 VSNs (no local UNI present on Boundary Node)	9,600
Maximum number of IP multicast S,Gs when operating as a boundary node (inter-area)	4,800
DvR host routes redistributed across area boundary	13,900
SPBM multicast-FIB entries	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	VSP 4900 Series	4,000	4,000
	VSP 7400 Series	4,000	4,000
6	VSP 4900 Series	3,500	4,000
	VSP 7400 Series	3,500	4,000
10	VSP 4900 Series	2,900	4,000
	VSP 7400 Series	2,900	4,000
20	VSP 4900 Series	2,000	4,000
	VSP 7400 Series	2,000	4,000
48	VSP 4900 Series	1,000	2,000
	VSP 7400 Series	1,000	2,000
72	VSP 4900 Series	750	1,500
	VSP 7400 Series	750	1,500
100	VSP 4900 Series	550	1,100
	VSP 7400 Series	550	1,100
128	VSP 4900 Series	450	900
	VSP 7400 Series	450	900
250	VSP 4900 Series	240	480
	VSP 7400 Series	240	480
<p>Note: Expect longer boot times with high scaled adjacency environments.</p>			

Interoperability Considerations for IS-IS External Metric

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

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

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

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.



Important Notices

[ExtremeCloud IQ Support](#) on page 62

[Compatibility with ExtremeCloud IQ Site Engine](#) on page 62

[Feature-Based Licensing](#) on page 62

[Memory Usage](#) on page 63

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 [VOSS 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 and Premier licenses. Premier licenses enable advanced features not available in the Base License.

For more information about licensing including feature inclusion, order codes, and how to load a license file, see [VOSS 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 64

[Restrictions and Expected Behaviors](#) on page 84

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 webservice Auth only for subject and root. For intermediate, it must be set with other required bits to avoid this issue. Add the root, intermediate CAs in the trust store of the browser for accessing the EDM with HTTPS.
VOSS-1265	On the port that is removed from a T-UNI LACP MLT, non T-UNI configuration is blocked as a result of T-UNI consistency checks.	When a port is removed from a T-UNI LACP MLT, the LACP key of the port must be set to default.
VOSS-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 .1p map, packets cannot follow correct internal QoS queues for FE tunnel to FE tunnel, or FE tunnel to regular NNI traffic.	Do not change the default ingress and egress .1p maps when using Fabric Extend. With default ingress and egress .1p maps, packets follow the correct internal QoS when using the Fabric Extend feature.
VOSS-1473	If the I-SID associated with a Switched UNI or Fabric Attach port does not have a platform VLAN association and you disable Layer 2 Trusted, then the non IP traffic coming from that port does not take the port QoS and still uses the .1p priority in the packet.	None.
VOSS-1530	If you improperly close an SSH session, the session structure information does not clear and the client can stop functioning.	Disable and enable SSH.

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 MHSa on a port, and create and configure RAVs for Non-EAP clients on the same port without DHCP Snooping and ND-inspection, no error is shown. The client displays as authenticated into RAV, even when port is not a member of RAV.	None.
VOSS-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 VOSS 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 3VSN 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>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.
VOSS-29220	In a scaled Multi-area SPB topology, after an event like an NNI link down, the failover time for multicast traffic can take up to 23-25 seconds when using 5720 Series switches as the boundary node pair.	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.
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.	Manually configure Fabric Extend tunnels between these models and ExtremeCloud SD-WAN.

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 [VOSS User Guide](#).

General Restrictions and Expected Behaviors

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

Table 27: 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 27: 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	OVSDB remote VTEP and MAC details can take between 5 to 10 minutes to populate and display after a HW-VTEP reboots.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.

Table 27: 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-12395	<p>You cannot use the following cables on 10 Gb fiber interfaces, or 40 Gb channelized interfaces, with the QSA28 adapter:</p> <ul style="list-style-type: none"> • 1, 3, and 5 meter QSFP28 25 Gb DAC • 20 meter QSFP28 25 Gb AOC 	n/a

Table 27: General restrictions (continued)

Issue number	Description	Workaround
VOSS-17871	Starting with VOSS 8.1.5, internal system updates have resulted in a more accurate accounting of memory utilization. This can result in a higher baseline memory utilization reported although actual memory usage is not impacted.	Update any network management alarms that are triggered by value with the new baseline.
VOSS-18523	When you configure a port using Zero Touch Provisioning Plus (ZTP+) with ExtremeCloud IQ Site Engine, the port cannot be part of both a tagged VLAN and an untagged VLAN.	n/a
VOSS-18774	SSL negotiation fails when using OpenSSL client version 1.1.1. With OpenSSL 1.1.1, the server-name extension is used. This extension needs to equal the domain name in the server certificate, otherwise the certificate lookup on the server fails because the FIPS 140-2 certified cryptographic module processes the server-name extension.	Can connect using: bash# openssl s_client -connect <domain-name>:443
VOSS-18851	Do not define a static route in which the NextHop definition uses an Inter-VRF redistributed route. Such a definition would require the system to perform a double lookup. When you attempt to define a static route in this way, an error message is generated.	Define the static route in such a way that it does not require Inter-VRF redistributed routing.
VOSS-21620	When interior nodes are running software earlier than Release 8.4 and a Multi-area takeover occurs between the boundary nodes (when the non-designated boundary node transitions to designated) in the network, the interior nodes might detect a false duplicate case between the stale LSP of the old virtual node and the new virtual node. This has no functional impact in the network.	n/a

Table 27: General restrictions (continued)

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

Table 27: General restrictions (continued)

Issue number	Description	Workaround
wi01142142	When a multicast sender moves from one port to another within the same BEB or from one vIST peer BEB to another, with the old port operationally up, the source port information in the output of the show ip igmp sender command is not updated with new sender port information.	<p>You can perform one of the following workarounds:</p> <ul style="list-style-type: none"> On an IGMP snoop-enabled interface, you can flush IGMP sender records. <p>Caution: Flushing sender records can cause a transient traffic loss.</p> <ul style="list-style-type: none"> On an IGMP-enabled Layer 3 interface, you can toggle the IGMP state. <p>Caution: Expect traffic loss until IGMP records are built after toggling the IGMP state.</p>
wi01145099	IP multicast packets with a time-to-live (TTL) equal to 1 are not switched across the SPB cloud over a Layer 2 VSN. They are dropped by the ingress BEB.	To prevent IP multicast packets from being dropped, configure multicast senders to send traffic with TTL greater than 1.
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	<p>When you disable EAPoL globally:</p> <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.

Table 27: General restrictions (continued)

Issue number	Description	Workaround
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
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 27: 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.

Redirect Next-hop Filter Restrictions

This feature does not behave the same way on all platforms:

On VSP 7400 Series, the redirect next-hop filter redirects packets with a time-to-live (TTL) of 1 rather than sending them to the CPU where the CPU would generate ICMP TTL expired messages. IP Traceroute does not correctly report the hop. For more information, see [VOSS User Guide](#).

Filter Restrictions

The following table identifies known restrictions.

Table 28: 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 on VSP 7400 Series.
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.
VSP 7400 Series	VLAN ID and VLAN_DOT1p attributes for untagged traffic are not supported for ingress/egress filters.
All platforms	Scaling numbers are reduced for IPv6 filters.

Table 28: ACL restrictions (continued)

Applies To	Restriction
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.
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 29: ACE restrictions

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



Resolved Issues this Release

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

- VOSS 8.10.3

Issue number	Description
CFD-10097	IP addresses not reachable after upgrade to 8.10.1.0 due to ARP entries being stuck on TX-NNI.
CFD-10165 CFD-10245 CFD-10467	5320 and 5420 switches can reboot randomly with high link utilization rate on MACsec-enabled links.
CFD-10359	rate-limit broadcast on NNI is lost after reboot.
VOSS-29245	Using ZTP and Auto-Sense to onboard a new switch dual-homed to a BCB layer, which has IS-IS keys preconfigured, the new switch does not receive DHCP replies from the uplink switches.
VOSS-29454	ARP replies are not received on a side IP interface configured between v1ST peers.
VOSS-29459	Connectivity issues to various devices until ARP cleared for destination VLAN.
VOSS-29470	Routed traffic failure from wireless users after roaming from Cisco Wireless solution side to Extreme Cloud wireless side.
VOSS-29540	Intermittent Layer 2 VSN traffic issues can occur if you remove and recreate the same C-VLAN I-SIDs simultaneously on multiple SPBM-enabled nodes.
VOSS-29647	If you disable and reenable LLDP on a port with an FA interface, the FA interface is recreated but without the authentication key.
VOSS-29684	tTacacspTask crash when using a password with more than 127 characters.
VOSS-29823	DHCP Server subnet EDM procedures, in the User Guide and EDM Help, include steps to select the subnet, and then select an Extra button. The Extra button does not appear in EDM.
VOSS-29839	The device name for 5320-16P-4XE-DC displays as 5320-16P-4XE in CLI show command output, ExtremeCloud IQ, and ExtremeCloud IQ Site Engine.



Related Information

MIB Changes on page 95

MIB Changes

Deprecated MIBs

Table 30: 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 31: Common

Object Name	Object OID	Modified in Release	Modification
rcCliRWAPassword	1.3.6.1.4.1.2272.1.19.2	8.10.1	CHANGE_RANGE: Changed range from 0..32 to 0..80
rcCliRWAPassword	1.3.6.1.4.1.2272.1.19.4	8.10.1	CHANGE_RANGE: Changed range from 0..32 to 0..80
rcCliRWL3Password	1.3.6.1.4.1.2272.1.19.6	8.10.1	CHANGE_RANGE: Changed range from 0..32 to 0..80

Table 31: Common (continued)

Object Name	Object OID	Modified in Release	Modification
rcCliRWL2Password	1.3.6.1.4.1.2272.1.19.8	8.10.1	CHANGE_RANGE: Changed range from 0..32 to 0..80
rcCliROPassword	1.3.6.1.4.1.2272.1.19.10	8.10.1	CHANGE_RANGE: Changed range from 0..32 to 0..80
rcCliRWL1Password	1.3.6.1.4.1.2272.1.19.16	8.10.1	CHANGE_RANGE: Changed range from 0..32 to 0..80
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)
avFabricAttachDiscElementsElementTypes	1.3.6.1.4.1.45.5.46.1.11.1.2	9.0	ADD_ENUM: faRing(18)
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)
rcSisLogicalInterfaceSrcIPAddr	1.3.6.1.4.1.2272.1.63.26.1.31	9.0	OTHER: Updated description to be available on all platforms

Table 31: Common (continued)

Object Name	Object OID	Modified in Release	Modification
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: fabricEngine5520x24TACDCACDC, fabricEngine5520x48TACDC, fabricEngine5520x48SEACDC, fabricEngine5520x24XACDC, fabricEngine752048YE8CE
rcVossSystemMgmtPortLedStatus	1.3.6.1.4.1.2272.1.101.1.1.1	9.0	OTHER: Update description to include 7520-48YE-8CE
rcVlanMvpngsidStatus	1.3.6.1.4.1.2272.1.3.2.1.84	9.0.2	ADD_NEW_VALUE: not-configured(3) OTHER: Updated description

Table 31: Common (continued)

Object Name	Object OID	Modified in Release	Modification
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

Table 32: VSP 4900 Series

Object Name	Object OID	Modified in Release	Modification
rcVirtualServiceScalarsName	1.3.6.1.4.1.2272.1.101.1.1.12.7	8.6	OTHER: Add rcVirtualServiceFigwCli in description
rcIsisLogicalInterfaceNextHopVrf	1.3.6.1.4.1.2272.1.63.26.1.13	8.8	Replaced read-only with read-create. Description changed.
bspePethPsePortPowerClassifications	1.3.6.1.4.1.45.5.8.1.1.15	8.10	OTHER: Updated description to include 5720 platform

Table 33: VSP 7400 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)
rcIsisGlobalMAHomeAlwaysUp	1.3.6.1.4.1.2272.1.63.1.33	8.6	OTHER: Changed DEFVAL from "false" to true"
rcVirtualServiceScalarsName	1.3.6.1.4.1.2272.1.101.1.1.12.7	8.6	OTHER: Add rcVirtualServiceFigwCli in description
rcIsisLogicalInterfaceNextHopVrf	1.3.6.1.4.1.2272.1.63.26.1.13	8.8	Replaced read-only with read-create. Description changed.

New MIBs

Table 34: Common

Object Name	Object OID	New in VOSS Release
rcnIsisPlsbSdwanTunnelStateTrap	1.3.6.1.4.1.2272.1.21.0.366	8.10.1
rcIsisPlsbLogicalIntfld	1.3.6.1.4.1.2272.1.63.9.46	8.10.1

Table 34: Common (continued)

Object Name	Object OID	New in VOSS Release
rcIsisPlsbLogicalIntfDestIPAddr	1.3.6.1.4.1.2272.1.63.9.47	8.10.1
rcIsisPlsbLogicalIntfSrcIPAddr	1.3.6.1.4.1.2272.1.63.9.48	8.10.1
rcIsisPlsbLogicalIntfSdwanState	1.3.6.1.4.1.2272.1.63.9.49	8.10.1
rcIsisLogicalInterfaceSdwanState	1.3.6.1.4.1.2272.1.63.26.1.34	8.10.1
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
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

Table 34: Common (continued)

Object Name	Object OID	New in VOSS Release
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

Table 35: VSP 4900 Series

Object Name	Object OID	New in Release
rcVossSystemAutoVimSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.8	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
rcIlsPlsbNickNameOrigin	1.3.6.1.4.1.2272.1.63.4.1.19	8.10
rcIlsPlsbNickNameServerSysId	1.3.6.1.4.1.2272.1.63.4.1.20	8.10
rcIlsPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	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 35: VSP 4900 Series (continued)

Object Name	Object OID	New in Release
rcChasPowerSupplyDetailVoltageOutput	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
rcChasPowerSupplyDetailCurrentOutput	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
rcChasPowerSupplyDetailPowerOutput	1.3.6.1.4.1.2272.1.4.8.2.1.21	9.0.2

Table 36: VSP 7400 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
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

Table 36: VSP 7400 Series (continued)

Object Name	Object OID	New in Release
rcIsisPlsbNickNameServerHostName	1.3.6.1.4.1.2272.1.63.4.1.21	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

Obsolete MIBs

Table 37: 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