VOSS 9.4 Release Notes

New Features, Improvements, and Known Issues

9039505-00 Rev AB
April 2026

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/

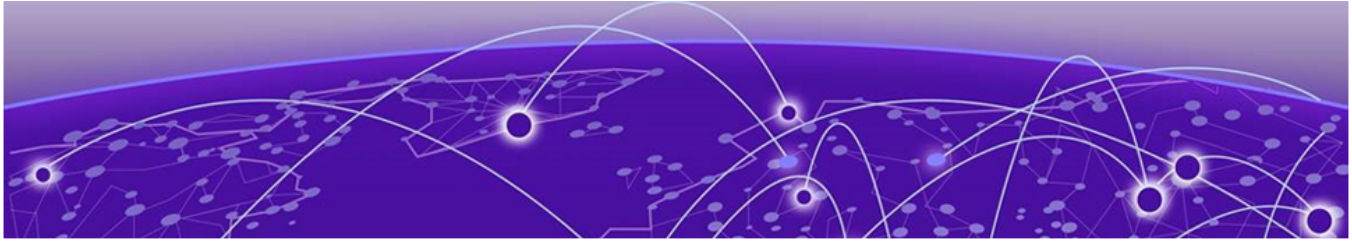

Table of Contents

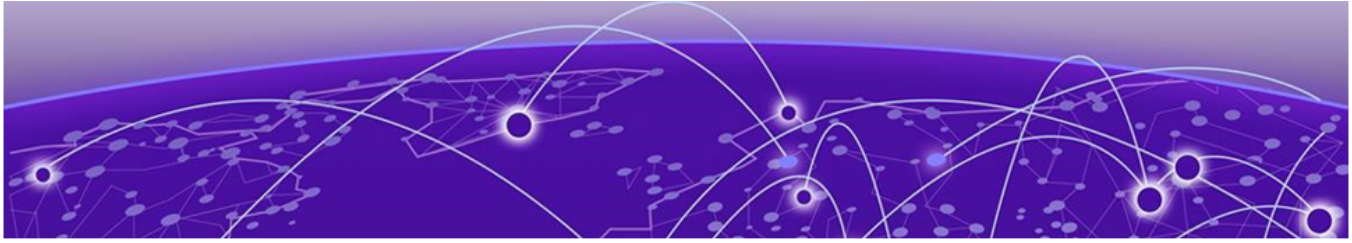

Abstract

The release notes for Extreme Networks VOSS version 9.4 detail new software features, upgrade considerations, scaling data, known issues, and resolved defects for the ExtremeSwitching VSP 4900 and VSP 7400 Series platforms. Software enhancements span Fabric, operational, platform, and security domains, including Segmented VRF for traffic isolation across trust levels, Auto-sense Link Debounce for PXE device support, IPv6 discard static routes, MSTP Restricted Role and TCN, PTPv2 Transparent Clock with VLAN support, and Enhanced Secure Mode hardening for TLS, SSH, and SSL ciphers. TPVM is updated to Ubuntu 24.04. Scaling information is updated for Segmented VRF impact. Targeted at network engineers and administrators with advanced knowledge of SPB Fabric and enterprise switching infrastructure.

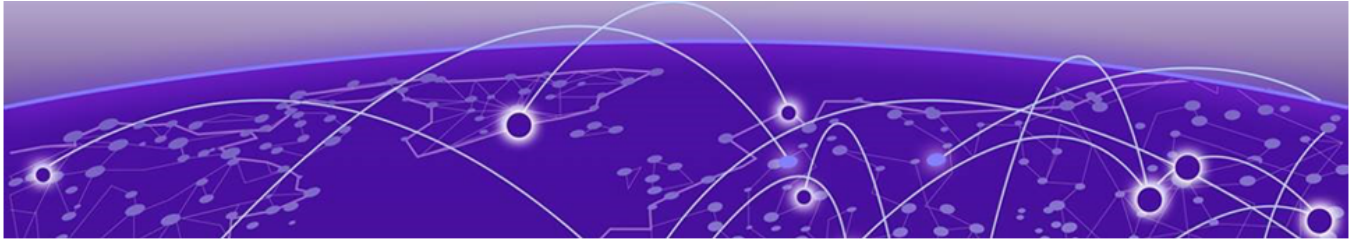

Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

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 *enter* and *type*	When you see the word *enter* in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says *type*.
Key names	Key names are written in boldface, for example **Ctrl** or **Esc**. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press **Ctrl**+**Alt**+**Del**
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 `cfm maintenance-domain maintenance-level <0-7>` , you can enter `cfm maintenance-domain maintenance-level 4`.
Bold text	Bold text indicates the GUI object name you must act upon. Examples: • 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 `ip address {A.B.C.D}`, 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 `show clock [detail]`, you can enter either `show clock` or `show clock detail`.
Ellipses (...)	An ellipsis (...) indicates that you repeat the last element of the command as needed. For example, if the command syntax is `ethernet/2/1 [ <parameter> <value> ]...`, you enter `ethernet/2/1` and as many parameter-value pairs as you need.
Italic Text	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: • `show ip route` • `Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]`

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 `access-policy by-mac action { allow | deny }`, you enter either `access-policy by-mac action allow` or `access-policy by-mac action deny`, 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 Product-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: 9.4 Release Notes revision changes

Revision	Change
AA	Initial revision for new release, see New in this Release on page 14
AB	Updated New Software Features or Enhancements on page 14

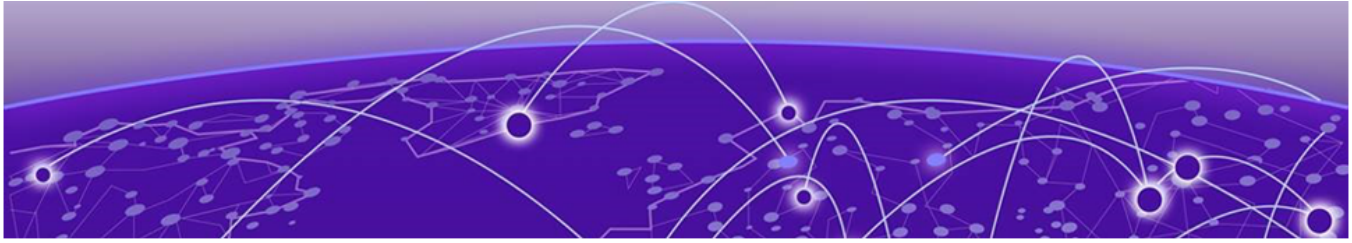

New in this Release

The following platforms support VOSS 9.4:

- ExtremeSwitching VSP 4900 Series
- ExtremeSwitching VSP 7400 Series

Note

For a specific list of supported models in each switch series, see Hardware and Software Compatibility on page 34.

For MIB-related changes, see MIB Changes on page 91.

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.

Hardware

VSP 4900 Series VIMs

VIM5-2Y and VIM5-4Y are end of service life and no longer supported. References to these VIMs are removed from the documentation. For information about these VIMs, see documentation for releases earlier than 9.4. For support information, see VSP 4900 Series Hardware on page 34.

New Software Features or Enhancements

The following sections describe what is new in this release.

Fabric Enhancements

The software supports the following Fabric enhancements:

- Auto-sense guest I-SID control—In this release, you can now disable the use of the onboarding I-SID as a guest I-SID.

  Note
  This functionality applies to Auto-sense ports in UNI-ONBOARDING, FA, and VOICE states. It does not apply to FA-PROXY, FA-PROXY-NOAUTH, FA-PROXY-RING, or NNI states.

- Auto-sense Link Debounce— You can use Auto-sense to configure Link Debounce on all Auto-sense UNI ports or only those that connect to Preboot Execution Environment (PXE) devices (by configuring it to *auto*). This enhancement addresses the following issue: PXE devices make a port bounce after the initial DHCP Discovery. This port bounce restarts the Auto-sensestate machine, the port exits the UNI state, the MAC is de-authenticated, and the wait-timeout period restarts. This process repeats in an endless loop, preventing the PXE device from receiving and applying the configuration. After the port bounce, the PXE device sends three DHCP Discovery messages during the first 6 seconds. The Link Debounce *auto* configuration enables the Link Debounce option after a well-known PXE packet is recognized on a port, which avoids a link bounce and allows the DHCP discovery packets to pass through successfully and ensure proper onboarding of PXE booted devices on Auto-sense ports.

  Note
  With this introduction in 9.4, you can no longer use the interface-level **link-debounce** command on Auto-sense ports. Use the new **auto-sense link-debounce** command instead.

- Auto-sense port state logging—This release creates log messages after an Auto-sense port transitions to a new state, using the following format: `0x0000c626 - 00000000 GlobalRouter HW INFO Auto-sense port <slot/port> entered [xyz] state`. The following message is an example of the port transitioning to the UNI state: `GlobalRouter HW INFO Auto-Sense port 1/5 entered UNI state`.
- Auto-sense SD-WAN VRF configuration enhancements—This release introduces the following enhancements to SD-WAN dynamic VRF and local breakout (LBO) BGP configuration:
  - You can now convert a dynamic SD-WAN VRF to a static VRF, which is necessary for SD-WAN high availability (HA) branch deployments.
  - The BGP configuration on the SD-WAN assigned VRF (or GRT) now includes a BGP origin field to track whether the configuration is dynamically created by Auto-sense or is statically configured.
  - You can now convert the dynamic LBO BGP configuration on the SD-WAN VRF to a static configuration, which is necessary for SD-WAN HA branch deployments.
  - The SD-WAN local breakout configuration on a user VRF now persists even when the SD-WAN Auto-sense port goes down, which is necessary for SD-WAN HA branch deployments.

- You can use the **ip bgp restart-bgp vrf sd-wan** command if the SD-WAN VRF origin or BGP origin is AUTO-SENSE.
- IPv6 discard routes—This release supports IPv6 discard static routes. A discard static route is a route with an invalid next hop that results in traffic matching the route (and not matching the more specific routes) being discarded.
- Segmented VRF (and thus Segmented Layer 3 VSN)—Use a Segmented VRF to control, isolate, and secure traffic flows across the network by creating isolated routing domains within a single, traditional VRF or Layer 3 VSN. The switch separates the VRF into three access areas, or trust levels: trusted, unrestricted, and untrusted. The switch classifies the traffic into those segments by the ingress VLAN. An untrusted VLAN, and thus I-SID, can only reach unrestricted segments. Clients on trusted VLANs can reach trusted VLANs as well as unrestricted VLANs. Typically, position a firewall on an unrestricted VLAN, which means it can respond to untrusted as well as trusted VLANs. Position unsecure IoT devices on untrusted VLANs, which means they are restricted to only reach a subset (unrestricted) destinations.

  Devices on an untrusted VLAN or I-SID can communicate with each other and can form an isolated communication zone or group that can only reach unrestricted (firewall) destinations. If devices must only talk to unrestricted (firewall) destinations, use a Private VLAN (PVLAN) to further isolate the devices from each other and direct them all to an unrestricted firewall interface.

  You can use a Segmented VRF to segment customers, internal services, functions, or security zones.

  

  Note

  This feature requires a Premier license if used with Layer 3 VSN.

- **show isis lsdb detail** now displays Layer 2 VSN mapping bit for TLVs 185 and 186.

For more information, see *VOSS User Guide* and *VOSS Command References*.

Operational Enhancements

The software supports the following operational enhancements:

- Autotopology—Starting in 9.4, when a device boots in Zero Touch Deployment (ZTD) mode, SONMP (also known as autotopology) is disabled by default. The first **save config** after ZTD boot includes **no autotopology** in the running configuration. Existing running configurations are unaffected. The CLI defaults, MIBs, and manual configuration options remain unchanged.
- CDP and LLDP on the same port—You can now configure a port to send and receive both CDP and LLDP packets. In previous releases, you could not enable CDP and LLDP on the same port.
- DHCP Option 150 for DHCP Server—Configure a list of up to eight TFTP server IP addresses. Clients, such as IP Phones, can request Option 150 from a DHCP Server to obtain configuration files or other information from a TFTP server.

- EDM changes—This release introduces the following changes:
  - EDM now displays clear status notifications for both successful and failed save operations.
  - The EDM navigation pane now includes a button to close all open tabs.
  - EDM access level passwords now support up to 80 characters. In previous releases, access level passwords could not exceed 32 characters.
  - The **MACsec KA Key** tab located in **Configuration** > **Security** > **Data Path** > **MACsec** now includes the **KeyStatus** field, which provides MKA Key status for Valid, Expired, and In-Use.
  - **PortMembers** fields are added to the **Fabric** > **IS-IS** > **Protocol Summary** tab to show MLT port members for IS-IS Interfaces and IS-IS Adjacency View.
  - Security and Qos Group fields located in **Security** > **Data Path** > **Advanced Filters (ACE/ACLs) ACL** tab are now renamed Primary and Secondary banks.
  - The **Security** > **Control Path** > **General** > **Web** tab adds a **InUseCertType** field.
  - The **SrcVrfID** parameter on **IP** > **<Protocol>** > **Redistribute** tabs was read-only in previous release. You can now configure this value.
  - VLAN lists now provide check boxes to select multiple VLANs simultaneously. Previously, you had to press and hold the **Ctrl** key when selecting multiple VLANs.
- In previous releases, when you configure the date on the switch, the maximum configurable year was 2038. In this release, support extends to year 2100. This change is implemented in CLI, EDM, and SNMP.
- IPFIX configuration per port—On IPFIX-supporting platforms, you can enable or disable IPFIX on all NNI ports or on individual UNI ports.
- Logging of **`show khi performance-scaling`** watermarks reached—After a resource monitored and displayed in the **`show khi resource-scaling`** command output reaches 80%, 90%, and 100%, the switch logs a WARNING message, sends an SNMP trap, and sets an alarm. If you use Extreme Cloud management applications, the switch also sends a message to that application.
- Multiple Spanning Tree Protocol Restricted Role (Root Guard) and Restricted TCN—This release introduces MSTP Restricted Role and Restricted TCNs. MSTP Restricted role prevents a port from accepting superior BPDUs from non-root bridges. When triggered, it puts the port in a root-inconsistent state. Restricted TCN prevents a switch port from propagating topology-change messages to other ports. When enabled, the port blocks and ignores all received TCNs.
- The **`quick-config-mgmt`** CLI command now supports the **Tab** key for command autocompletion. The command must be complete to run it.
- **`show fulltech`** command—This command now includes output from the **`show khi resource-scaling`** and **`show io resources`** commands.
- **`show io l2-tables`** command—This command output now includes EEPROM data for ports with an inserted optical pluggable component.
- **`show khi resource-scaling`** command—This command output now provides a clearer view of how switches allocate and share hardware resources.
- TFTP Block Number Rollover—In this release, the software uses block number rollover functionality to transfer files larger than 32 MB using TFTP. With this functionality, when the block number reaches 65,535, it resets to 0 while it maintains

the standard block size of 512 bytes and allows the transfer to continue seamlessly. In earlier releases, TFTP data blocks were numbered sequentially, which caused the transfer to stop when the block number reached 65,535 and limited the maximum file size.

- TTL handling of bridged traffic with routed SPB IP Multicast—This release adds Layer 2 VSN-based forwarding to IP Multicast over Fabric Connect. When a multicast sender and receiver reside in the same Layer 2 VSN, the Fabric bridges the multicast traffic instead of routing it. The switch preserves the IP TTL and the source C-MAC address, which prevents TTL-related packet drops for low-TTL protocols such as PTPv1. The ingress BEB advertises TLV 188 to identify the sender and Layer 2 VSN, and receiving BEBs use this information to select bridged or routed forwarding for each multicast stream. This feature operates across all Multi-area Fabric Connect topologies.
- VRF name autocompletion—You can use the CLI command completion features for VRF names in show and configuration commands.

For more information, see *VOSS User Guide* and *VOSS Command References*.

Platform Enhancements

The software supports the following platform enhancements:

- Domain resolution test for a specific DNS—When you enable dynamic IP configuration, the Network Service Probe interface uses the DNS server information it receives from the DHCP server to test DNS resolution. The interface obtains its IP address, default gateway, and up to three DHCP-advertised DNS servers. You can query any of these DNS servers—primary, secondary, or tertiary—to verify connectivity.
- Hardware Watchdog Reset—The Hardware Watchdog monitors a communication heartbeat that the software transmits to the hardware subsystem. If this heartbeat is interrupted, or not detected, a system-level hardware reset is hardware-initiated to restore the device to a known-good operational state. This feature is a critical fail-safe mechanism to ensure system reliability and availability.
- Secure log file transfer—You can now configure log file transfer, with the **`logging transferFile {1-10}`** command, to use Secure Copy (SCP) rather than TFTP or FTP.
- Third Party Virtual Machine (TPVM) OS update—The version of Linux in the TPVM image is updated to Ubuntu 24.04. A new image file is available in 9.4. For more information, see File Names for this Release on page 20.
- VLAN-based Transparent Clock for PTPv2 —This release introduces Precision Time Protocol version 2 (PTPv2) Transparent Clock with VLANs. This feature improves time-synchronization accuracy across the network by compensating for switch latency in PTP timing messages. The switch measures the residence time of PTP packets and updates the Correction Field as packets traverse the network, which maintains accurate end-to-end timing between the timeTransmitter Clock and timeReceiver Clock within a VLAN.

This feature requires a Premier license.

- VSP 7400 Series minimum fan speed—Use the new **`sys fan set-min-speed <20-100>`** command to configure the minimum fan speed as a percentage between the minimum supported speed and the maximum supported speed and adjust the thermal operations of the switch. This value is the minimum speed at which the fan operates; the switch increases the fan speed when necessary.

For more information, see *VOSS User Guide* and *VOSS Command References*.

Security Enhancements

The software supports the following security enhancements:

- Enhanced Secure Mode (ESM)—If the switch operates in Enhanced Secure Mode, this release introduces the following changes:

  Note
  The switch generates audit logs if you try to enable unapproved algorithms or key exchange methods in ESM.

  - The default minimum TLS version for Syslog is TLS 1.2.
  - RSA SHA224 and ECDSA SHA224, SHA256, SHA384, and SHA512 are disabled during the SSL handshake.
  - For the web server, DHE ciphers are disabled during SSL handshake.
  - The following SSH key exchange methods are disabled by default:
    - diffie-hellman-group-exchange-sha256
    - diffie-hellman-group14-sha1
  - The following additional SSH encryption types are disabled by default:
    - aes192-cbc
    - aes192-ctr
    - rijndael128-cbc
    - rijndael192-cbc
  - SSH host key algorithm x509v3-ssh-rsa is no longer allowed.
- SSH packet size—The maximum SSH packet size is 35840 bytes.

For more information, see *VOSS User Guide* and *VOSS Command References*.

Inclusion of 9.3.1

This release includes the following 9.3.1 feature changes:

- OpenAPI Enhancements:
  - **`openapi local-mgmt ttl <60-86400>`** command—Use this command to configure the time-to-live (ttl) value for the authentication token.
  - **`openapi local-mgmt minimum-tls {1.2 | 1.3}`** command—Use this command to configure the minimum TLS version.

- **show application openapi log** command—This command now includes the *reverse* parameter to display Open API logs entries in chronological order.

- RADIUS VSA Enhancements:
  - Additional parameters for Extreme-Dynamic-Client-Assignments Vendor Specific Attribute (VSA) used in RADIUS for dynamic VLAN and PVLAN assignment:
    - `none`—Use an existing VLAN or PVLAN instead of creating a new one.
    - `igmpqaddr=<IPv4 address>`—Configure the IGMP Querier address for traffic within the VLAN either for IGMP Snooping or for Multicast Lite or Routed Multicast.

    The updated string format to create a dynamic VLAN is as follows:

    ```
    create=vlan|pvlan|none, pv=Primary VLANID, sv=secondary VLANID,
    vni=L2-ISID, ev=EGRESS-VLAN-tag, vn=vlan-name, vnin=isid-name,
    mvni=MVPN-ISID, igmpqaddr=<IPv4 address>
    ```
  - Additional IGMP features for Extreme-Dynamic-Config Vendor Specific Attribute (VSA) used in RADIUS:
    - IGMP version 3 (IGMPV3)
    - IGMP Fast Leave (IGMPFAST)
- **show license** command—This command output is enhanced with extra information to explicitly indicate when Premier features are included with Extreme Platform ONE Networking licenses.
- SSH to IS-IS system-ID—If you perform a factory reset on a switch, it also resets the SSH key. Forming new SSH connections to this switch would fail because the remote host had changed. The only recourse was to also factory reset the switch originating the SSH connection. Starting with 9.3.1, you can resolve this situation by deleting the known hosts file on the originating switch with the **delete /intflash/.ssh/known_hosts** CLI command .

Other Changes

New File

File Names for this Release on page 20 includes a new file for the RADIUS dictionary.

Scaling Updates

VRF Scaling on page 65 is updated for Segmented VRF impact.

File Names for this Release

Important

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

After you download the software, calculate and verify the md5 checksum. For more information, see *VOSS User Guide*.

When extracting the software image file, the extraction process appends the software version portion of the extracted file names to include the final full software version. (For example, extracting **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 file names and sizes for this release.

Table 5: VSP 4900 Series Software File names and Sizes

Description	File	Size
RADIUS dictionary	dictionary.fabricengine	3,502 bytes
Fabric IPsec Gateway	FabricIPSecGW_VM_5.2.0.0.ova	4,034,211,840 bytes
YANG model	restconf_yang.tgz	506,020 bytes
Third Party Virtual Machine (TPVM)	TPVM_Ubuntu24.04_04_01April 2026.qcow2	3,381,665,280 bytes
Logs reference	VOSS4900.9.4.0.0_edoc.tar	40,058,880 bytes
MD5 Checksum files	VOSS4900.9.4.0.0.md5	671 bytes
MIB - supported object names	VOSS4900.9.4.0.0_mib_sup.txt	1,586,270 bytes
MIB - objects in the OID compile order	VOSS4900.9.4.0.0_mib.txt	8,731,837 bytes
MIB - zip file of all MIBs	VOSS4900.9.4.0.0_mib.zip	1,293,989 bytes
Open source software - Master copyright file	VOSS4900.9.4.0.0_oss-notice.html	2,597,473 bytes
SHA512 Checksum files	VOSS4900.9.4.0.0.sha512	1,878 bytes
Software image	VOSS4900.9.4.0.0.tgz	327,215,246 bytes
EDM Help files	VOSSv9.4.0_HELP_EDM_gzip.zip	5,281,220 bytes

Table 6: VSP 7400 Series Software File names and Sizes

Description	File	Size
RADIUS dictionary	dictionary.fabricengine	3,502 bytes
Fabric IPsec Gateway	FabricIPSecGW_VM_5.2.0.0.ova	4,034,211,840 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 6: VSP 7400 Series Software File names and Sizes (continued)

Description	File	Size
Third Party Virtual Machine (TPVM)	TPVM_Ubuntu24.04_04_01April 2026.qcow2	3,381,665,280 bytes
Logs reference	VOSS7400.9.4.0.0_edoc.tar	40,058,880 bytes
MD5 Checksum files	VOSS7400.9.4.0.0.md5	671 bytes
MIB - supported object names	VOSS7400.9.4.0.0_mib_sup.txt	1,588,532 bytes
MIB - objects in the OID compile order	VOSS7400.9.4.0.0_mib.txt	8,731,837 bytes
MIB - zip file of all MIBs	VOSS7400.9.4.0.0_mib.zip	1,293,989 bytes
Open source software - Master copyright file	VOSS7400.9.4.0.0_oss-notice.html	2,597,473 bytes
SHA512 Checksum files	VOSS7400.9.4.0.0.sha512	1,878 bytes
Software image	VOSS7400.9.4.0.0.tgz	328,365,326 bytes
EDM Help files	VOSSv9.4.0_HELP_EDM_gzip.zip	5,281,220 bytes

Upgrade and Downgrade Considerations

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:

- DHCP Server vendor options configuration change
- 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
- 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 25
- Switches That Will Use Zero Touch Deployment on page 25

Impact of Auto-sense Port Configuration in Release 9.3

Important

In Release 9.3 and later, if an Auto-sense port on a switch without an IS-IS Hello Authentication key connects to an Auto-sense port on another switch with an IS-IS Hello Authentication key, both ports transition to the NNI-AUTH-FAIL state and dynamically enable STP multi-homing. If you onboard one or more access switches, ensure all core switches that receive the access switch uplinks run Release 9.3 or later.

Failure to run 9.3 or later on the core switches can cause Spanning Tree loops on the onboarding VLAN between those switches. For more information about the NNI-AUTH-FAIL state, see Auto-sense Port States in *VOSS User Guide*.

IS-IS Route Tagging

Caution

To use IS-IS Route Tagging on GRT IS-IS routes, you must also configure the metric-type as external. If you want to use IS-IS tags on GRT as internal routes, all Fabric nodes must be above a minimum software version. Any switch in the SPB Fabric that runs earlier software versions triggers an exception if you use metric type internal. To ensure this does not occur, if you attempt to configure a tag and the metric-type is not external, the switch reminds you to upgrade the software on all devices. You must ensure all devices in the network run the minimum required software.

Table 7: Minimum software required

NOS	Minimum software versions
Fabric Engine	8.10.6.1 and later 9.0.5.1 and later 9.1 and later
VOSS	8.10.6.1 and later 9.0.5.1 and later 9.1 and later
VSP 8600 Series	8.1.7 and later

Validated Upgrade Paths

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

Note

For any versions prior to 8.10.0.0 or 9.2.0.0, an intermediate upgrade is recommended because pre-8.10.0.0 and pre-9.2.0.0 versions are not validated.

For non-validated upgrade paths, perform the upgrade with one or two switches initially before doing a widespread upgrade.

Table 8: Validated upgrade paths

Product	8.10.x to 9.4	9.2.x to 9.4	9.3.x to 9.4
VSP 4900 Series	Y	Y	Y
VSP 7400 Series	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 29 and *VOSS User Guide*.
2. Upgrade to this release from one of the previously described releases, see Validated Upgrade Paths on page 24.
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 24.
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.
- The Auto-sense guest I-SID is disabled 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 File Names for this Release on page 20. 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 28.

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 24.

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

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 24.

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.

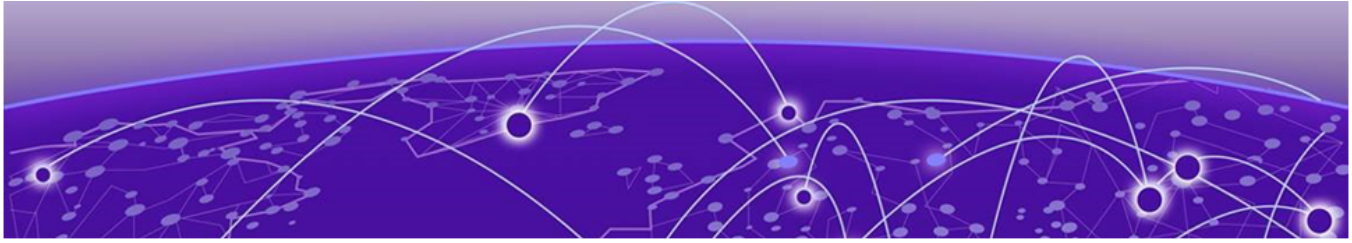

Hardware and Software Compatibility

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

VSP 4900 Series Hardware

Table 9: Switch models

Model	Initial release	Supported new VOSS feature release				
		9.0.3	9.1	9.2	9.3	9.4
VSP4900-48P	8.1	Y	Y	Y	Y	Y
VSP4900-12MXU-12XE	8.1.5	Y	Y	Y	Y	Y
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				
		9.0.3	9.1	9.2	9.3	9.4
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	N
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	N

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	VIM5-4YE	VIM5-2Q
Number of supported ports for VSP4900-48P and VSP4900-24S	4	4	2	1
Number of supported ports for VSP4900-24XE and VSP4900-12MXU-12XE	4	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)	40 Gbps 10 Gbps (with channelization
PHY present	No	Yes	Yes	No
Copper transceiver support (1 Gbps/10 Gbps)	10GBASE-T only	Both	10GBASE-T only	Not applicable
MACsec	Not supported	128/256 bit	128/256 bit	Not supported
Forward Error Correction (FEC)	Not supported	Not supported	Default is Auto-FEC - FEC Auto, CL108, CL91, CL74 and No FEC supported	Not supported
1 Gbps Auto-Negotiation	Disabled	Enabled	Not applicable	Not applicable

Table 11: VSP 4900 Series VIM Matrix (continued)

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

VSP 7400 Series Hardware

Part number	Model Number	Initial release	Supported new VOSS feature release				
			9.0.3	9.1	9.2	9.3	9.4
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 the 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, see *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, see *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)

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*.

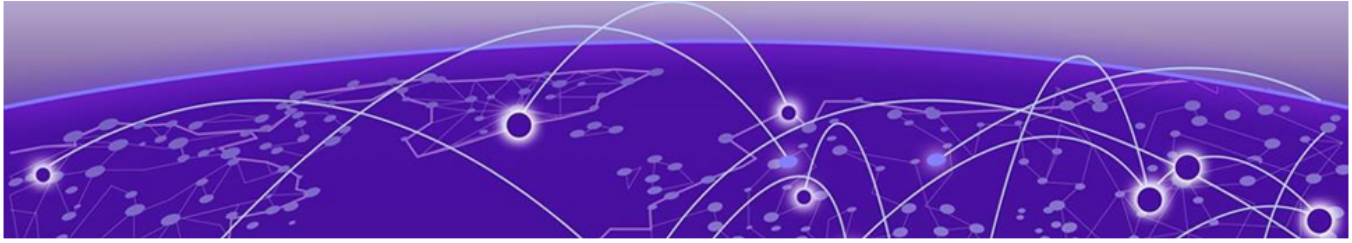

Scaling

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	120,000
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 42.
	VSP 7400 Series	200 See Maximum Number of Directed Broadcast Interfaces on page 42.
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 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: • 2,000 for VSP4900-48P with RESTCONF • 1,000 for VSP4900-24S with RESTCONF	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

Table 12: Layer 2 Maximums (continued)

Attribute	Product	Maximum number supported
SLPP VLANs	VSP 4900 Series	128
	VSP 7400 Series	500 VLANs with a minimum SLPP counter of 0.5 seconds 1,000 VLANs with a minimum SLPP counter of 1 second 2,000 VLANs with a minimum SLPP counter of 2 seconds
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 43.
	VSP 7400 Series	200 See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 43.

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 assuming there is only 1 NLB cluster IP interface per 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.

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. 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.

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 47.
	VSP 7400 Series	1,000 See IP Interface Maximums for VSP 7400 Series on page 47.
VRRP interfaces (IPv4 or IPv6) **Note:** Do not create more than 10 IPv6 VRRP VRs on a single VLAN.	VSP 4900 Series	252 See IP Interface Maximums for VSP 4900 Series on page 47.
	VSP 7400 Series	500 per switch 256 per VRF See IP Interface Maximums for VSP 7400 Series on page 47.

Table 13: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
Anycast IP Gateway interfaces	VSP 4900 Series	250 See IP Interface Maximums for VSP 4900 Series on page 47.
	VSP 7400 Series	500 250 on boundary node See IP Interface Maximums for VSP 7400 Series on page 47.
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 47.
	VSP 7400 Series	499 See IP Interface Maximums for VSP 7400 Series on page 47.
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
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

Table 13: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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 65.	VSP 4900 Series	256 including mgmt VRF and GRT See IP Interface Maximums for VSP 4900 Series on page 47.
	VSP 7400 Series	256 including mgmt VRF and GRT See IP Interface Maximums for VSP 7400 Series on page 47.
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
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
DHCP client addresses provided by the DHCP server	VSP 4900 Series	10,000 clients
	VSP 7400 Series	100,000 clients
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

Table 13: IP Unicast Maximums (continued)

Attribute	Product	Maximum number supported
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
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
Virtual router IDs with Anycast IP Gateway	VSP 4900 Series	16
	VSP 7400 Series	16

IP Interface Maximums Clarification

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

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

IP Interface Maximums for 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) + (#Anycast IP Gateway VLANs if Anycast IP Gateway router)
- 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 + (#Anycast IP Gateway VLANs if Anycast IP Gateway router)

For additional detail, see IP Interface Maximums Clarification on page 47.

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) + (#Anycast Gw VLANs if Anycast Gw router) 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) 2x(#Anycast Gw VLANs if Anycast Gw router) 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) + (#Anycast Gw VLANs if Anycast Gw router) 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) 2x(#Anycast Gw VLANs if Anycast Gw router) cannot exceed 1000

For additional detail, see IP Interface Maximums Clarification on page 47.

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 48.
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
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

Table 17: IP Multicast Maximums (continued)

Attribute	Product	Maximum number supported
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
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 43.

Table 18: DvR Maximums

Attribute	Product	Maximum number supported
Note: • 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. Note: A DvR domain containing only Controller nodes and no Leaf nodes can have more than 8 Controllers per domain.	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 53.		
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
Maximum number of IP Source Guard filters	VSP 4900 Series	100
	VSP 7400 Series	48-port model: 480 32-port model: 320
Total IPv4 Egress rules/ACEs (Port based, Security filters)	VSP 4900 Series	248
	VSP 7400 Series	783 271 if you enable **boot config flags ipv6-egress-filter**
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:
    - (num ACLs + num security ACEs) <= 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:
    - (num ACLs + num ACEs) <= 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 (security) - # of ACLs) + (512 (QoS) - # 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 - # 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:
    - num ACLs <= 512 && (num ACLs + num Primary ACEs) <= 767 && (num ACLs + num Secondary ACEs) <= (767 – X) where X = num IPv6 ACLs + 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:
  - num IPv6 ACLs <= 383 && (num IPv6 ACLs + num ACEs) <= (767 – X) where X = num non-IPv6 ACLs + 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:
      - num ACLs <= 254 && (num ACLs + num Security ACEs) <= 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:
    - num ACLs <= 256 && (num ACLs + num Security ACEs) <= 512

This maximum implies a single port on outPort ACLs.

The switch supports the following maximum limits for 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 - num non-IPv6 egress ACLs
- 511 IPv6 egress ACEs

  This theoretical maximum implies
  - 1 egress ACL with 511 Security ACEs
  - a port member count of 1 for outPort ACLs
  - 511 - num IPv6 egress ACLs

Routed Private VLANs/E-TREEs Scaling

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

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

Table 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 Note: These resources are shared with the IPv4 Filter Ingress rules/ACEs.	VSP 4900 Series	382 hosts
	VSP 7400 Series	767 hosts

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 65.	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 62
	VSP 7400 Series	See Maximum Number of SPB Multicast Data I-SIDs on page 62

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 62 **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 62 **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 SD-WAN tunnels signaled on an Auto-sense port	VSP 4900 Series	115
	VSP 7400 Series	125

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 3VSNs/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
Number of nodes that can function as Multi-area SPB boundary nodes between two areas	4 in a non-vIST configuration, 2 in a vIST configuration
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

Table 26: Multi-area SPB maximums (continued)

Scaling	VSP 7400 Series
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

Note:
Expect longer boot times with high scaled adjacency environments.

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 l1-hellointerval`** and **`isis l1-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 l1-hellomultiplier**, instead of using the default value of 3.

VRF Scaling

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

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.

Segmented VRF Impact on Scaling

Segmented VRFs use more system resources than traditional VRFs (Layer 3 VSNs). A Segmented VRF consumes three system resources rather than one system resource like a traditional VRF.

The following example illustrates this impact by showing a switch with two configured VRFs, one traditional VRF (vrf69) and one Segmented VRF (vrf99). The **show ip vrf** output indicates four VRFs with two VRF names and the output of the **show khi resource-scaling** command indicates four Layer 3 VSN resources are used. This count of four is because vrf99 consumes three resources.

```
Switch:1#show ip vrf
=============================================================================================================================================================
                                                          VRF INFORMATION
=============================================================================================================================================================
VRF       VLAN      ARP       RIP       OSPF      BGP       PIM       NBRv6     RIPng     OSPFv3    PIM6      VRF IDs
COUNT     COUNT     COUNT     COUNT     COUNT     COUNT     COUNT     COUNT     COUNT     COUNT     COUNT     ALLOCATED
-------------------------------------------------------------------------------------------------------------------------------------------------------------
-
4         7         13        1         1         1         1         0         1         1         1         6

VRF                  VRF  VLAN   ARP                                  NBRv6                       UNICAST  SD-WAN LOCAL  VRF
NAME                  ID   COUNT  COUNT  RIP    OSPF   BGP    PIM    COUNT  RIPng  OSPFv3 PIM6  ACTIVE   BREAKOUT      ORIGIN    SEGMENTED
-------------------------------------------------------------------------------------------------------------------------------------------------------------
-
GlobalRouter          0    4      3      TRUE   TRUE   TRUE   TRUE   0      TRUE   TRUE   TRUE  TRUE     FALSE         DYNAMIC   FALSE
vrf99                 1    0      0      FALSE  FALSE  FALSE  FALSE  0      FALSE  FALSE  FALSE TRUE     FALSE         CONFIG    TRUE
vrf69                 69   1      7      FALSE  FALSE  FALSE  FALSE  0      FALSE  FALSE  FALSE TRUE     FALSE         CONFIG    FALSE
MgmtRouter            512  1      0      FALSE  FALSE  FALSE  FALSE  0      FALSE  FALSE  FALSE TRUE     FALSE         DYNAMIC   FALSE

4 out of 4 Total Num of VRF Entries displayed.
Switch:1#show khi resource-scaling
=============================================================================================================
                                        KHI resource-scaling
=============================================================================================================
Item                                    Maximum       Maximum     Currently  Available   Usage   Shared resource
                                        (theoretical) (actual)    used                   (%)     group
------------------------------------------ ------------- ----------- ---------- ----------- ------- ------------------
##<snip>##
Services
-------------------------------------------------------------------------------------------------------------
Multicast-fib entries                   10000         10000       8          9992          1 %
L2VSNs                                  500           500         4          496           1 %   RES1
Transparent UNIs                        29            29          0          29            0 %   RES1,RES3
Switched UNIs                           400           400         0          400           0 %   RES1,RES3,RES3
Private VLANs                           100           100         1          99            1 %
L3VSNs                                  64            64          4          60            6 %   RES1,RES3
```

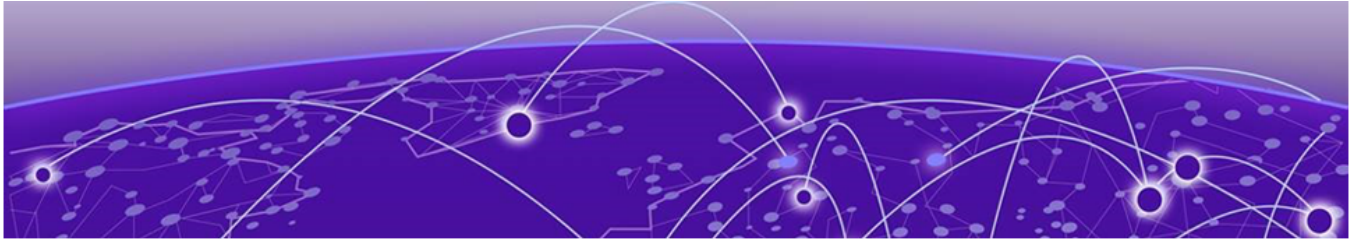

Important Notices

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

Platform Overview and Integration Updates

This section outlines the capabilities, integrations, and version-specific updates across the following Extreme Networks core platforms.

ExtremeCloud™ IQ

ExtremeCloud IQ is a cloud-managed networking solution that delivers unified, full-stack management for wireless access points, switches, and routers. It supports:

- ◦ Device onboarding and configuration
  ◦ Real-time monitoring and troubleshooting
  ◦ Advanced reporting and analytics

Leveraging machine learning and artificial intelligence, ExtremeCloud IQ processes millions of data points—from the network edge to the data center—to generate actionable insights and enable intelligent automation across the network.

Switches running VOSS support zero touch connection to ExtremeCloud IQ. Zero touch deployment with ExtremeCloud IQ simplifies and accelerates device provisioning and configuration.

VOSS 9.4 was successfully tested with ExtremeCloud IQ version 25.10.

The switch software integrates with ExtremeCloud IQ using IQAgent.

ExtremeCloud IQ Site Engine

Zero Touch Provisioning Plus (ZTP+) enables you to deploy and configure switches in ExtremeCloud IQ Site Engine with minimal server configuration and intervention.

VOSS 9.4 was successfully tested with ExtremeCloud IQ Site Engine version 26.5.10.

Extreme Platform ONE Networking

Extreme Platform ONE Networking provides a unified foundation for integrating various Extreme applications, including ExtremeCloud IQ, Extreme Platform ONE Security, ExtremeCloud SD-WAN, and Extreme Intuitive Insights into a single, AI-powered platform that simplifies network deployment, management, and security.

VOSS 9.4 was successfully tested with Extreme Platform ONE Networking version 25.10.0.

Licensing

The switches support a perpetual licensing model that includes Base, Premier, and Premier with MACsec licenses. Premier and Premier with MACsec licenses enable use of advanced features not available in the Base License. To see which features a platform supports, see *Fabric Engine and VOSS Feature Matrix*.

For more information about licensing including feature inclusion, order codes, and how to load a perpetual license file on the switch, see *VOSS User Guide*.

Management CLIP Preferred for Management Client Applications

If you configure both a Management VLAN in routing mode and a Management CLIP, the switch prefers the Management CLIP when it selects the source IP address and outgoing interface for traffic initiated by management client applications, for example RADIUS reachability. For management client traffic, ensure connectivity exists in the Management CLIP context.

Memory Usage

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

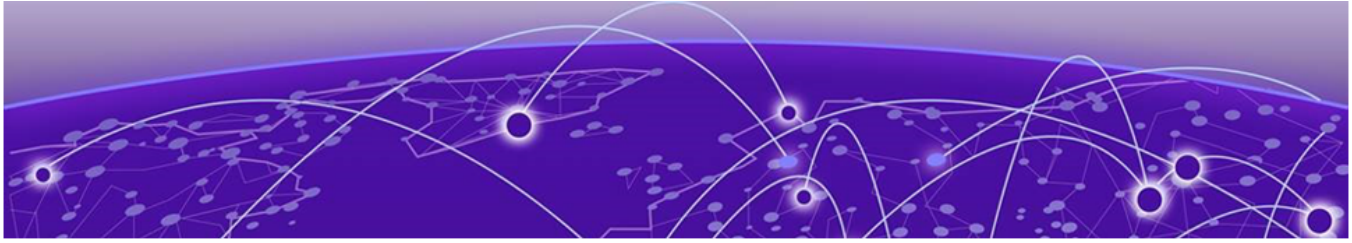

Known Issues and Restrictions

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
CFD-14948	A data forwarding issue was identified in environments utilizing vIST where MAC addresses move rapidly between access points (APs). This behavior resulted in a mismatch between hardware and software forwarding tables, leading to packet drops.	None
CFD-15355	Wireless clients on FA dynamic VLANs fail to get DHCP IP address when no platform VLAN exists for the same. A fix for this issue is targeted for 9.5.0.0 release.	None
CFD-16205	An unexpected reboot may occur when frequent LLDP state transitions cause memory exhaustion. A fix for this issue is targeted for 9.3.3.0 release.	None
CFD-16262 CFD-16555	Ping to IS-IS source IP may not work. A fix for this issue is targeted for 9.3.3.0 release.	None
CFD-16483 VOSS-34911	MHSA authentication on port should override all previous states and assignments A fix for this issue is targeted for 9.3.3.0 release.	None

Issue number	Description	Workaround
	HTTPS connection fails for CA-signed certificate with certificate inadequate type error on FF.	Ensure End-Entity, Intermediate CA and Root CA certificates are all SHA256 based and RSA2048 key signed, and Extended key usage field is set to TLS webserver Auth only for subject and root. For intermediate, it must be set with other required bits to avoid this issue. Add the root, intermediate CAs in the trust store of the browser for accessing the EDM with HTTPS.
VOSS-1285	CAKs are not cleared after setting the device to factory-default.	None. Currently this is the default behavior and does not affect functionality of the MACsec feature.
VOSS-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-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-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.

Issue number	Description	Workaround
VOSS-10815	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. 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.	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.
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: `http(s)://<ip-address>:8080/apps/restconfdoc/`. For more information, see *VOSS User Guide*.
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-15541	You can experience temporary traffic loss when shutting down an LACP SMLT port (and therefore causing the local SMLT to go down), in a network with scaled Multicast traffic over an SPB cloud, while the datapath processes all dpm letter messages during LCAP recovery. This slow LACP recovery situation is only seen with scaled Multicast traffic over an SPB cloud.	Use static MLTs.
VOSS-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, andVSP4900-12MXU-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-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.

Issue number	Description	Workaround
VOSS-20455	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: • `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-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-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: • 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: • 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 *`e1000`* 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.
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.

Issue number	Description	Workaround
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-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-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.
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.

Issue number	Description	Workaround

VOSS-28625	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:
	`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`
	The issue is caused by the interoperability of this specific ACL configured to mirror the I-SID traffic, and the Multi-area filters.
		Remove the ACL used to mirror I-SID traffic on the boundary node. Use Fabric RSPAN (Mirror to I-SID) to achieve similar functionality.
		Alternatively, use matchtype "uniOnly" instead of “both".

VOSS-28672	IPFIX does not learn MCoSPB NNI-UNI flows on 7520 Series, 7720 Series, and VSP 7400 Series.	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-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-30287	An intermittent connectivity issue occurs over a Fabric Extend destination tunnel in a failover scenario when the IS-IS unicast FIB computation does not point to the shortest path.	This situation is temporary. You can perform an action, such as configuring any same I-SID on the Fabric Extend tunnel ends to trigger an IS-IS computation. Wait for the IS-IS computation to generate.
VOSS-30292	If IPv6 Shortcuts are explicitly disabled, SSH connections does not work on VSP 4900 Series.	Enable IPv6 Shortcuts.
VOSS-30576	`WARNING: CPU: 0 PID: 1 at kernel/rcu/tree_plugin.h:297 rcu_note_context_switch+0x44/0x340` kernel message, which displays on the console during bootup has no functional impact on 4220 Series and 5320 SeriesXT. .	None.
VOSS-30980	After you enable an IP VPN instance on a VRF that you configure for the IS-IS logical interface and the adjacency establishes through the Fabric Extend tunnel with a nickname server, Dynamic Nickname Offers are discarded on the port with the Fabric Extend tunnel.	As a workaround, disable and then reenable IP VPN on the VRF or use the automatic nickname.

Issue number	Description	Workaround
VOSS-30990	When you change the advanced-fabric-bandwidth-reservation flag from low to high, you cannot enable Auto-sense on ports reserved as loopback ports after you reboot the switch. An example of the message that displays is as follows: `Cleanup for auto-sense failed on port: 1/10, reason: VLAN cleanup failed!`	As a workaround, disable Auto-sense on reserved loopback ports before you reboot the switch.
VOSS-31315	The following message displays when you upgrade to VOSS Release 9.1 on VSP 4900 Series:`DMAR: [Firmware Bug]: No firmware reserved region can cover this RMRR [0x000000003e2e0000-0x000000003e2fffff], contact BIOS vendor for fixes.`	None.
VOSS-31352	When you disconnect a Fabric Attach client from an Auto-sense port and reconnect a device that does not transmit LLDP packets, the device displays the wrong port default VLAN ID when the port transitions from the WAIT to UNI state.	As a workaround, bring the port down and then bring the port up to restart the Auto-sense state on the switch to display the correct default VLAN ID when the port transitions to the UNI state.
VOSS-31465	When an IPv6 RSMLT in the forwarding state cannot ping the IPv6 link-local address of the vIST peer in that particular RSMLT VLAN, local routes whose next-hop is the link-local address of the vIST peer can fail.	None.
VOSS-32147	ZTP+ fails to discover and connect to ExtremeCloud IQ Site Engine, in a scenario where two DNS servers are configured on the switch through DHCP but, although it is running, the primary DNS server cannot resolve the extremecontrol hostname.	In the DHCP server configuration, the primary DNS server must be able to resolve "extremecontrol". If one of the DNS servers cannot resolve the extremecontrol hostname, for example, in the case of a public DNS server, add that server in the DHCP configuration with a lower priority.
VOSS-32270	When ARP entries exceed the 8000 limit on VLANs without an assigned I-SID on 5520 Series, multiple error messages display and ARP entries fail to program correctly when you add additional ARP entries.	As a workaround, assign I-SIDs to VLANs that can manage more than 8000 ARP entries.
VOSS-32312	The following message displays on the console: `unable to rotate the file '/intflash/shared/telegraf.log', rename /intflash/shared/telegraf.log /intflash/shared/telegraf.<timestamp>.log: no such file or directory.`	None. You can safely ignore this message.

Issue number	Description	Workaround
VOSS-32476	In a scenario with multiple misconfigurations in single and multiple areas such as Multi-area SPB inter-area duplicate nickname/system-ID recovery, log messages can be mismatched.	None.
VOSS-33191	On the 7830 Series, traffic packets smaller than 64 bytes are drop but the `TOO SHORT` counter does not increment as expected.	None.
VOSS-33478	When you insert a 40G break-out cable but you do not channelize it, partner devices can still detect a valid signal and bring the links up. As a result, all 10G lanes on the break-out cable show link-up status, even though the 40G interface is not channelized.	As a workaround, channelize the 40G break-out cable.
VOSS-33685	In a scenario that uses AUTO-MLT with IS-IS backup adjacency, the MLT interface remains even after you disable all member ports. Although the IS-IS adjacency correctly goes down when all ports in the MLT are shut down, the MLT itself remains empty.	None.
VOSS-33697	During early system startup, an IQ Agent core file can generate due to a race condition where the Redis context is not yet available when accessed by the IQ Agent. This crash is highly intermittent and harmless, as the system lifecycle management automatically restarts the IQ Agent, allowing normal operation to resume without any service impact.	None.
VOSS-33821	An intermittent, timing-dependent issue can occur when IS-IS on interior nodes and multi-area (IS-IS remote) on boundary nodes are bounced at the same time. In certain cases, the non-designated boundary node can take significantly longer than expecte, up to 10 minutes, to become fully multi-area operational. During this period, only the designated boundary node forwards traffic.	Bounce any adjacency on the designated boundary node within the home area.
VOSS-33970	After a 5320-16P-2MXT-2X reboot, the following message can display: `Card did not respond to voltage select! : -110 / do_gpt: mmc dev 0 NOT available`, and the switch does not boot to CLI.	Physical access is required to manually power-cycle the system.
VOSS-34052	On 5320 Series, 5420 Series, 5520 Series, and 5720 Series, after you enable PTP Transparent Clock, all IPFIX flows on all ports are exported with timestamp fields set to 0.	Disable PTP Transparent Clock, save the configuration, and reboot the switch.
VOSS-34101	7830 Series: After de-channelizing a 400 Gbps port that was configured as 4x25G or 4x10G, the port can begin to flap continuously. The link does not recover on its own.	Perform a manual reset of the affected port to restore link functionality.
VOSS-34648	On 7830 Series, the **`show ip ipfix flows`** command does not display individual learned IPFIX flows. Use the **`show ip ipfix`** command to see the total number of learned flows.	None.

Issue number	Description	Workaround
VOSS-34859	When both Anycast IP Gateway and an IPv6 interface are configured on the same VLAN, bouncing IS-IS can disrupt IPv6 traffic. As a result, IPv6 traffic on that VLAN does not recover.	Delete and recreate the IPv6 interface.
VOSS-34898	After configuring an IPv6 recursive static route in the GlobalRouter VRF—where the configured next hop is reachable across the SPB cloud and resolves to an internal IPv6 special/hidden IPv6 Shortcut address—the route becomes active as expected. However, after saving the configuration and rebooting the switch, the IPv6 recursive route does not return to an active state. The route remains inactive until its status is manually bounced.	Take either of the following actions: • Disable and re-enable the affected IPv6 recursive static route. • Delete and recreate the IPv6 recursive static route.
VOSS-34968	Port does not go operationally down when DDM alarms are raised and **`pluggable-optical-module ddm-alarm-portdown ddm-monitor`** is configured.	None.

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.

Table 27: General restrictions (continued)

Issue number	Description	Workaround
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-2166	The IPsec security association (SA) configuration has a NULL Encryption option under the **Encrpt-algo** parameter. Currently, you must fill the **encrptKey** and **keyLength** sub-parameters to set this option; however, these values are not used for actual IPsec processing as it is a NULL encryption option. The NULL option is required to interoperate with other vendors whose IPsec solution only supports that mode for encryption.	There is no functional impact due to this configuration and it only leads to an unnecessary configuration step. No workaround required.
VOSS-21946	When you create a vrf using the POSTMAN API platform, special characters, such as \\\\ and ### included in the URL are ignored.	None.
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.

Table 27: General restrictions (continued)

Issue number	Description	Workaround
VOSS-7640	The same route is learned via multiple IPv6 routing protocols (a combination of two of the following : RIPng, OSPFv3 and BGPv6). In this specific case, an eBGP (current best – preference 45) route is replaced by and iBGP (preference 175) which in turn is replaced by and OSPFv3 (external 2) route (preference 125).	None.
VOSS-7647	With peer group configuration, you cannot configure Update Source interface with IPv6 loopback address in EDM.	Use CLI.
VOSS-9174	OVSDB remote VTEP and MAC details can take between 5 to 10 minutes to populate and display after a HW-VTEP reboots.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.
VOSS-9462	OVSDB VNID I-SID MAC bindings are not populated on HW-VTEPs after configuration changes.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.
VOSS-10168	The system CLI does not prevent you from using the same IP address for the VXLAN Gateway hardware VTEP replication remote peer IP and OOB Management IP.	Manually check the IP configured as the OOB Management IP. Do not use the OOB Management IP address as the replication remote peer IP address.
VOSS-11817	The OVS connect-type for virtual service Vports is designed in such a way that it connects to any generic virtual machine (VM) guest OS version using readily available Ethernet device drivers. This design approach provides initial connectivity to the VM in a consistent manner. A consequence of this approach is that Vports created with connect-type OVS will show up as 1 Gbps interfaces in the VM even though the underlying Ethernet connection supports 10 Gbps .	If additional performance is desired, upgrade the VM guest OS with an Ethernet device driver that supports 10 Gbps interfaces.

Table 27: General restrictions (continued)

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-12151	If logical switch has only hardware ports binding, and not VM behind software VTEP, Broadcast, Unknown Unicast, and Multicast (BUM) traffic does not flow between host behind two hardware VTEP. The NSX replicator node handles the BUM traffic. NSX does not create the replicator node unless a VM is present. In an OVSDB topology, it is expected that at least one VM connects to the software VTEP. This issue is an NSX-imposed limitation.	After you connect the VM to the software VTEP, the issue is not seen.
VOSS-13794	You cannot use SFTP to transfer files larger than 2 GB to the switch.	Use SCP.
VOSS-15391	An SNMP walk on the **`rcIgmpSnoopTraceTable`** table will fail with an **`OID not increasing`** error. CLI and EDM are unaffected by this issue.	None.
VOSS-17871	Starting with VOSS 8.1.5, internal system updates have resulted in a more accurate accounting of memory utilization. This can result in a higher baseline memory utilization reported although actual memory usage is not impacted.	Update any network management alarms that are triggered by value with the new baseline.

Table 27: General restrictions (continued)

Issue number	Description	Workaround
VOSS-18238	When a management VLAN with DHCP is used to reach a RADIUS server, and the RADIUS server cannot be reached, the system waits for 15 minutes before attempting to reach the RADIUS server again. This is true even if the RADIUS server becomes reachable before the 15 minutes have elapsed.	None.
VOSS-18278	On the 5520 Series switch, when you make any change relating to port speed, the port statistics are cleared. This applies to all front panel fiber and copper ports as well as VIM ports. The following are examples of changes relating to port speed: • 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-18523	When you configure a port using Zero Touch Provisioning Plus (ZTP+) with ExtremeCloud IQ Site Engine, the port cannot be part of both a tagged VLAN and an untagged VLAN.	n/a
VOSS-18851	Do not define a static route in which the NextHop definition uses an Inter-VRF redistributed route. Such a definition would require the system to perform a double lookup. When you attempt to define a static route in this way, an error message is generated.	Define the static route in such a way that it does not require Inter-VRF redistributed routing.

Table 27: General restrictions (continued)

Issue number	Description	Workaround
VOSS-21620	When interior nodes are running software earlier than Release 8.4 and a Multi-area takeover occurs between the boundary nodes (when the non-designated boundary node transitions to designated) in the network, the interior nodes might detect a false duplicate case between the stale LSP of the old virtual node and the new virtual node. This has no functional impact in the network.	n/a
wi01068569	The system displays a warning message that routes will not inject until the apply command is issued after the enable command. The warning applies only after you enable redistribution, and not after you disable redistribution. For example: **`Switch:1(config)#isis apply redistribute direct vrf 2`**	n/a
wi01112491	IS-IS enabled ports cannot be added to an MLT. The current release does not support this configuration.	n/a
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, `snmp_comm.txt`, 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

Table 27: General restrictions (continued)

Issue number	Description	Workaround
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
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.	You can perform one of the following workarounds: • On an IGMP snoop-enabled interface, you can flush IGMP sender records. **Caution:** Flushing sender records can cause a transient traffic loss. • On an IGMP-enabled Layer 3 interface, you can toggle the IGMP state. **Caution:** Expect traffic loss until IGMP records are built after toggling the IGMP state.
wi01171670	Telnet packets get encrypted on MACsec-enabled ports.	None.
wi01210217	The command **`show eapol auth-stats`** displays `LAST-SRC-MAC` for NEAP sessions incorrectly.	n/a
wi01212034	When you disable EAPoL globally: • 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

Table 27: General restrictions (continued)

Issue number	Description	Workaround
wi01212247	BGP tends to have many routes. Frequent additions or deletions impact network connectivity. To prevent frequent additions or deletions, reflected routes are not withdrawn from client 2 even though they are withdrawn from client 1. Disabling route-reflection can create a black hole in the network.	Bounce the BGP protocol globally.
wi01212585	LED blinking in EDM is representative of, but not identical to, the actual LED blinking rates on the switch.	n/a
wi01213066 wi01213374	EAP and NEAP are not supported on brouter ports.	n/a
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: • Enable PIM on the edge. • Ensure that IST peers are either RP or DR but not both.

Table 27: General restrictions (continued)

Issue number	Description	Workaround
wi01224683 wi01224689	Additional link bounce can occur on 10 Gbps ports when toggling links or during cable re-insertion. Additional link bounce can occur with 40 Gbps optical cables and 40 Gbps break-out cables, when toggling links or during cable re-insertion.	n/a
wi01229417	Origination and termination of IPv6 6-in-4 tunnel is not supported on a node with vIST enabled.	None.
wi01232578	When SSH keyboard-interactive-auth mode is enabled, the server generates the password prompt to be displayed and sends it to the SSH client. The server always sends an expanded format of the IPv6 address. When SSH keyboard-interactive-auth mode is disabled and password-auth is enabled, the client itself generates the password prompt, and it displays the IPv6 address format used in the **ssh** command.	None.
VOSS-26218	In a scaled environment, running the **show io l2-tables** command reiteratively can cause the switch to reboot.	For scaled scenarios, do not run the **show io l2-tables** command in a loop.
VOSS-31214	With the upgrade to Mocana 7, RadSec Proxy certificates must conform to the following SP 800-132 specifications for PBKDFv2 parameters: • salt length of at least 128 bits • derived key length of at least 112 bits • iteration count of at least 1000 When you generate public or private key-pairs protected by a password with older versions of OpenSSL, the default PKCS5_SALT_LEN is 8 bytes, which results in TLS failures.	You can perform one of the following workarounds: • Recompile OpenSSL to use a PKCS5_SALT_LEN value of 16 bytes and generate new certificates. • Use a newer version of OpenSSL that is FIPS approved. • Generate the keys without password protection.

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.
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.

Table 29: ACE restrictions (continued)

Applies To	Restriction
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.

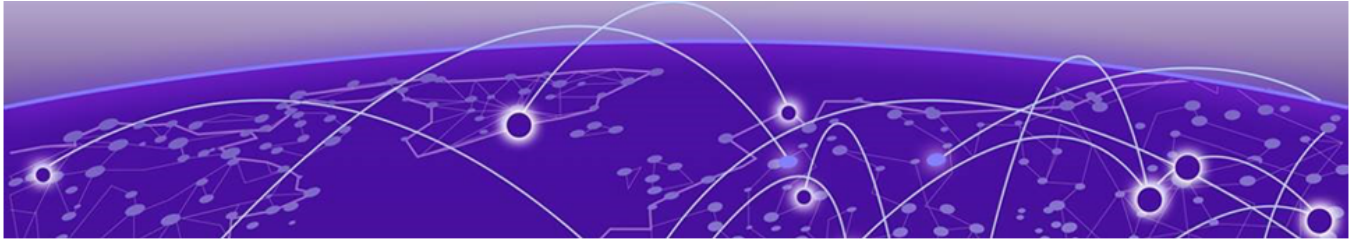

Resolved Issues this Release

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

- VOSS 9.1.3
- VOSS 9.2.3
- VOSS 9.3.2

Issue number	Description
CFD-12633	In rare cases on 5420F-16MW-32P-4XE, Power over Ethernet (PoE) can stop working following a software upgrade from Release 8.8.x and earlier. This issue is caused by a firmware update or initialization error during the upgrade process. `POE ERROR POE Board Initialization failure (unable to download firmware POE is OUT OF SERVICE)`
CFD-13522	HW INFO Sensor 14 in slot 1 overheat critical temperature alarm
CFD-13592	Multi-rate ports can stop working if frequent auto-negotiation changes on that port. This is specific to Flex-UNI ports when the I-SID they are part of has no platform VLAN associated. This issue can affect the first 16 ports of 5420F-16MW-32P-4XE or 5420M-16MW-32P-4YE.
CFD-13639	SPBM Multi-area - Prefixes installed wrongly on ISIS/OSPF ASBR routing table due to VN overload bit for best route election on interior nodes.
CFD-13825	In a scaled environment with a large number of OSPF routes, a race condition exists when you apply redistribution commands back-to-back for multiple source control protocols, which can result in an incorrect route reference counter. Subsequent loss of an impacted route results in traffic not rerouting properly.
CFD-14065	In environments using Equal-Cost Multi-Path (ECMP) routing for the default route (0.0.0.0/0), traffic can be incorrectly routed when a lower cost route replaces an existing one.

Issue number	Description
CFD-14656	Some devices, such as the POS terminals, send two GARP packets within the first two seconds after the port transitions to UP, and then remain silent. Node Alias is enabled in WAIT state; however, since the port is not a member of any VLAN, the packets are discarded. As a result, the devices' MAC and IP addresses do not reach the Auto-sense survive-reboot engine. The solution is to bypass Node Alias and forward the information directly to Auto-sense even when the port is not a member of any VLAN. The survive-reboot engine can then trigger PING and ARP requests to wake up the POS devices. Additionally, there are silent devices which, when pinged by the Auto-sense survive-reboot engine, respond with an ICMP Reply where the destination IP is set to the default gateway IP. VOSS and Fabric Engine simply bridge these packets instead of forwarding them to the CPU, preventing proper processing. The solution is to redirect ICMP traffic to the CPU whenever Node Alias is enabled on the port.
CFD-14884	VSP 7400 switch is locking up with ports down.
CFD-14993	5420: POS NEAP devices not authenticating on Auto-sense port.
CFD-15145	7520 - Some MACsec MKA links flap with `Communication Is Not Secure` messages
CFD-15423	EAP INFO Maximum allowed MAC reached, dropping MAC <MAC_Address> on Port <Port_No>.
CFD-15424	Loop occurs for few seconds before LACP SMLT comes up if the vIST is established on top of NNIs formed with Auto-sense.
CFD-15449	TDR test on specific SKUs of 5320, 5420, and 5520 caused traffic loss even after the test finished.
CFD-15532	Switch restart with core - port_state_fsm_task.
CFD-15919	Switch crashed after clearing the IS-IS LSDB.
CFD-15920	Default route leaked from GRT to Layer 3 VSN not working after change of Layer 3 VSN I-SID.
CFD-16092	When a nickname duplicate occurred within a single area, inter-area nickname-duplicate handling code was mistakenly run on the Boundary Nodes. This led to a crash when the nodes attempted to clean PLSB multicast entries in the other area.
CFD-16202	Switch crash with cbcp-main.x core dump and LACP handshake in the backtrace.
CFD-16531	High response time when querying larger MLT configuration using openAPI.
VOSS-32799	The 7830 Series does not display the hardware revision power supply units (PSUs).
VOSS-32873	In Release 8.10 and later, 5320 Series and 5420 Series do not support jumbo packet frames larger than 3748 bytes.

Issue number	Description
VOSS-33015	On the 7830 Series, if you configure the link speed at 10G on either the copper or SFP+ management port, the switch software detects the link speed as 1G.
VOSS-33615	In a scenario that involves inter-area system ID duplication, IP Shortcut traffic can fail to recover even after the duplicate system ID condition is resolved. ARP entries destined for an interior node can be incorrectly programmed on a tunnel pointing to a duplicate boundary node. These misprogrammed entries persist and prevent proper traffic forwarding, as they are not automatically corrected. This issue is caused by the RPF mechanism not functioning as expected for ARP packets, allowing entries to be programmed on incorrect tunnels.
VOSS-33808	On the 7830 Series, the EDM LED for the fiber management port incorrectly displays green when operating at 1G speed, instead of the expected amber blinking to indicate activity.
VOSS-33809	On the 7830 Series, the EDM LED for the fiber management port operating at 10G should display solid green for link-up and blinking green to indicate activity.
VOSS-33866	IP Anycast Gateway ONE-IP interfaces generate duplicate IP log reports when you enable router IS-IS. This occurs due to the IS-IS overload-on-startup flag, which acts as a hold-down timer for the Anycast Gateway. The default timer is 20 seconds. During this period, the ONE-IP interface cannot become operational. However, the VLAN IP interface becomes active and sends ARP requests using the VLAN chassis MAC address instead of the expected Anycast Gateway MAC.
VOSS-33872	When querying statistics on a front panel port on the 7830 Series, the `inDiscard` counter reports a higher number of dropped packets than expected. This behavior occurs because the counter includes link-local packets (such as LLDP, ISIS, LLC, and BPDU), and can also count locally processed packets per port. These packets are copied to the CPU and then dropped to prevent VLAN flooding.

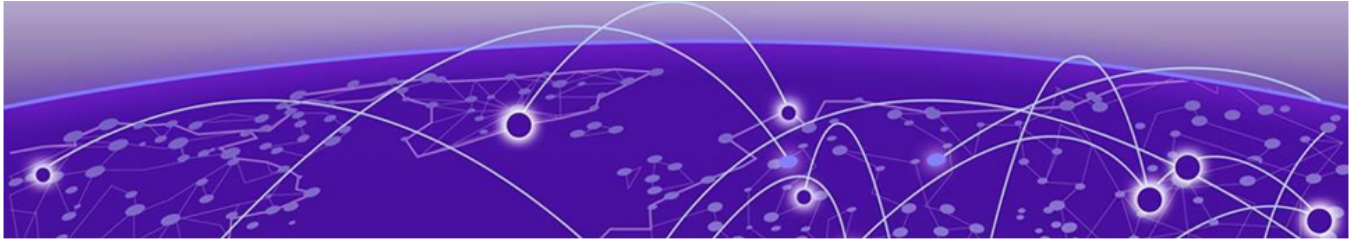

Related Information

MIB Changes on page 91

MIB Changes

Modified MIBs

Table 30: Common

Object Name	Object OID	Modified in Release	Modification
rcnIsisPlsbInterAreaDuplicateNicknameTrap	1.3.6.1.4.1.2272.1.21.0.368	9.3	changed rcIsisHomeSysId with rcIsisHomeChassisMac and rcIsisRemoteSysId with rcIsisRemoteChassisMac
rcMACSecIfCipherSuite	1.3.6.1.4.1.2272.1.88.2.1.5	9.3	OTHER: Changed default value from none(1) to gcmAes128(2)
bspePethPsePortPowerClassifications	1.3.6.1.4.1.45.5.8.1.1.1.15	9.3	OTHER: Update description to include all PoE platforms
rcWebRWAPassword	1.3.6.1.4.1.2272.1.18.3	9.3.1	CHANGE_RANGE: Changed the range from 1..32 to 1..80
rcWebROPassword	1.3.6.1.4.1.2272.1.18.7	9.3.1	CHANGE_RANGE: Changed the range from 1..32 to 1..80
rcWebMinimumPasswordLength	1.3.6.1.4.1.2272.1.18.32	9.3.1	CHANGE_RANGE: Changed the range from 1..32 to 1..80
rcUserSetTimeYear	1.3.6.1.4.1.2272.1.31.1	9.4	CHANGE_RANGE: Changed the range from 1998..2097 to 1998..2199
rcUserSetTimeYear	1.3.6.1.4.1.2272.1.31.1	9.4	CHANGE_RANGE: Changed the range from 1998..2199 to 1998..2100

Table 30: Common (continued)

Object Name	Object OID	Modified in Release	Modification
rc2kBootConfigEnableIpv6Mode	1.3.6.1.4.1.2272.1.100.5.1.47	9.4	OTHER: Update description for 7830 platform
rc2kCardSlotPower	1.3.6.1.4.1.2272.1.100.6.1.32	9.4	OTHER: Fixed typo in description. Changed "Administrately" to "Administratively"
rcDhcpServerGlobalTFTPServerIp	1.3.6.1.4.1.2272.1.232.1.1.1.5	9.4	OTHER: Update description with option 66
rcDhcpServerSubnetTFTPServerIp	1.3.6.1.4.1.2272.1.232.1.1.2.1.7	9.4	OTHER: Update description with option 66
rcnKhiAsicResourceUtilizationTrap	1.3.6.1.4.1.2272.1.21.0.370	9.4	Changed back rcKhiAsicResourceType to v9.3.0.0 and added new element "totalBmacs(27)"

Table 31: VSP 4900 Series

Object Name	Object OID	Modified in Release	Modification
rcIsisLogicalInterfaceNextHopVrf	1.3.6.1.4.1.2272.1.63.26.1.13	8.8	Replaced read-only with read-create. Description changed.
bspePethPsePortPowerClassifications	1.3.6.1.4.1.45.5.8.1.1.1.15	8.10	OTHER: Updated description to include 5720 platform

Table 32: VSP 7400 Series

Object Name	Object OID	Modified in Release	Modification
rcIsisLogicalInterfaceNextHopVrf	1.3.6.1.4.1.2272.1.63.26.1.13	8.8	Replaced read-only with read-create. Description changed.
rcPortType	1.3.6.1.4.1.2272.1.4.10.1.1.2	9.4	ADD_NEW_VALUE: rc100Gb40GbBiDi(274)

New MIBs

Table 33: Common

Object Name	Object OID	New in Release
rcMACSecKeychainAssociationId	1.3.6.1.4.1.2272.1.88.6.1.1	9.3
rcMACSecKeychainAssociationName	1.3.6.1.4.1.2272.1.88.6.1.2	9.3
rcMACSecKeychainAssociationRowStatus	1.3.6.1.4.1.2272.1.88.6.1.3	9.3
rcMACSecKeychainAssociationPortMembers	1.3.6.1.4.1.2272.1.88.6.1.4	9.3
rcMACSecKeychainAssociationKeyNum	1.3.6.1.4.1.2272.1.88.6.1.5	9.3
rcMACSecKeychainAssociationKeysCount	1.3.6.1.4.1.2272.1.88.6.1.6	9.3
rcMACSecKeychainId	1.3.6.1.4.1.2272.1.88.7.1.1	9.3
rcMACSecKeychainKeyId	1.3.6.1.4.1.2272.1.88.7.1.2	9.3
rcMACSecKeychainKeyCKN	1.3.6.1.4.1.2272.1.88.7.1.3	9.3
rcMACSecKeychainKeyCAK	1.3.6.1.4.1.2272.1.88.7.1.4	9.3
rcMACSecKeychainKeyExpiry	1.3.6.1.4.1.2272.1.88.7.1.5	9.3
rcMACSecKeychainKeyRowStatus	1.3.6.1.4.1.2272.1.88.7.1.6	9.3
rcMACSecIfKAName	1.3.6.1.4.1.2272.1.88.2.1.6	9.3
rcNlsMgmtVlanRouterMode	1.3.6.1.4.1.2272.1.223.1.1.19	9.3
rcNlsMgmtVlanRouterModeTable	1.3.6.1.4.1.2272.1.223.25	9.3
rcNlsMgmtVlanRouterModeEntry	1.3.6.1.4.1.2272.1.223.25.1	9.3
rcNlsMgmtVlanRouterModeVlanId	1.3.6.1.4.1.2272.1.223.25	9.3
rcNlsMgmtVlanRouterModeVrfName	1.3.6.1.4.1.2272.1.223.25.1.2	9.3
rcNlsMgmtVlanRouterModeRowStatus	1.3.6.1.4.1.2272.1.223.25.1.3	9.3
rcAutoSenseAutoMltEnable	1.3.6.1.4.1.2272.1.231.1.1.1.34	9.3
rcIsisCircuitPortMembers	1.3.6.1.4.1.2272.1.63.5.1.11	9.3
rcIsisAdjPortMembers	1.3.6.1.4.1.2272.1.63.10.1.6	9.3
rcAutoSenseFaProxyNoAuthForceAuth	1.3.6.1.4.1.2272.1.231.1.1.1.35	9.3
rcPortAutoSenseNoNni	1.3.6.1.4.1.2272.1.4.10.1.1.141	9.3
rcNlsMgmtConvertIpv6Address	1.3.6.1.4.1.2272.1.223.23.16	9.3
rcNlsMgmtConvertIpv6PrefixLength	1.3.6.1.4.1.2272.1.223.23.17	9.3
rcNlsMgmtConvertIpv6Gateway	1.3.6.1.4.1.2272.1.223.23.18	9.3
rcNlsMgmtConvertMode	1.3.6.1.4.1.2272.1.223.23.19	9.3
rcNlsMgmtConvertMoveDefaultStaticRoute	1.3.6.1.4.1.2272.1.223.23.20	9.3
rcRadiusServHostResolvedAddressType	1.3.6.1.4.1.2272.1.29.5.1.36	9.3
rcRadiusServHostResolvedAddress	1.3.6.1.4.1.2272.1.29.5.1.37	9.3
rcRadiusDynAuthClientResolvedAddressType	1.3.6.1.4.1.2272.1.29.6.1.8	9.3
rcRadiusDynAuthClientResolvedAddress	1.3.6.1.4.1.2272.1.29.6.1.9	9.3
rcnRateLimitExceededTrap	1.3.6.1.4.1.2272.1.21.0.369	9.3

Table 33: Common (continued)

Object Name	Object OID	New in Release
rcRateLimitAction	1.3.6.1.4.1.2272.1.14.25	9.3
rcRateLimitActionPollInterval	1.3.6.1.4.1.2272.1.14.25.1	9.3
rcRateLimitActionTable	1.3.6.1.4.1.2272.1.14.25.2	9.3
rcRateLimitActionEntry	1.3.6.1.4.1.2272.1.14.25.2.1	9.3
rcRateLimitActionIfIndex	1.3.6.1.4.1.2272.1.14.25.2.1.1	9.3
rcRateLimitIfActionShutdownTimeout	1.3.6.1.4.1.2272.1.14.25.2.1.2	9.3
rcRateLimitIfActionTrapInterval	1.3.6.1.4.1.2272.1.14.25.2.1.3	9.3
rcMltOrigin	1.3.6.1.4.1.2272.1.17.10.1.51	9.3
bspePethMainPoEDetectType	1.3.6.1.4.1.45.5.8.1.2.1.6	9.3
rcIgmpInterfaceExtnFastLeaveEnableOrigin	1.3.6.1.4.1.2272.1.30.1.1.43	9.3.1
rcIgmpInterfaceExtnExplicitHostTrackingEnableOrigin	1.3.6.1.4.1.2272.1.30.1.1.44	9.3.1
rcIgmpInterfaceExtnCompatibilityModeEnableOrigin	1.3.6.1.4.1.2272.1.30.1.1.45	9.3.1
rcIgmpInterfaceExtnVersionOrigin	1.3.6.1.4.1.2272.1.30.1.1.46	9.3.1
rcIgmpInterfaceExtnSnoopQuerierEnableOrigin	1.3.6.1.4.1.2272.1.30.1.1.47	9.3.1
rcIgmpInterfaceExtnSnoopQuerierAddrOrigin	1.3.6.1.4.1.2272.1.30.1.1.48	9.3.1
rcIgmpInterfaceExtnRoutedSpbQuerierAddrOrigin	1.3.6.1.4.1.2272.1.30.1.1.49	9.3.1
rcAutoSenseGuestIsidEnable	1.3.6.1.4.1.2272.1.231.1.1.1.36	9.4
rcAutoSenseLinkDebounceTimeout	1.3.6.1.4.1.2272.1.231.1.1.1.38	9.4
rcAutoSenseLinkDebounceType	1.3.6.1.4.1.2272.1.231.1.1.1.37	9.4
rcDhcpServerGlobalTftpServerEntry	1.3.6.1.4.1.2272.1.232.1.1.17.1	9.4
rcDhcpServerGlobalTftpServerOpt150Ip	1.3.6.1.4.1.2272.1.232.1.1.17.1.1	9.4
rcDhcpServerGlobalTftpServerTable	1.3.6.1.4.1.2272.1.232.1.1.17	9.4
rcDhcpServerGlobalTftpServerType	1.3.6.1.4.1.2272.1.232.1.1.17.1.2	9.4
rcDhcpServerSubnetTftpServerEntry	1.3.6.1.4.1.2272.1.232.1.1.18.1	9.4
rcDhcpServerSubnetTftpServerOpt150Ip	1.3.6.1.4.1.2272.1.232.1.1.18.1.3	9.4
rcDhcpServerSubnetTftpServerSubnetBitmask	1.3.6.1.4.1.2272.1.232.1.1.18.1.2	9.4
rcDhcpServerSubnetTftpServerSubnetIp	1.3.6.1.4.1.2272.1.232.1.1.18.1.1	9.4
rcDhcpServerSubnetTftpServerTable	1.3.6.1.4.1.2272.1.232.1.1.18	9.4
rcDhcpServerSubnetTftpServerType	1.3.6.1.4.1.2272.1.232.1.1.18.1.4	9.4
rcIpBgpGeneralGroupVrfBgpOrigin	1.3.6.1.4.1.2272.1.8.101.1.30	9.4
rcLldpPortCdpConfigDualModeAdminState	1.3.6.1.4.1.2272.1.220.1.2.1.1.3	9.4
rcNlsMgmtServiceProbeQueryDns	1.3.6.1.4.1.2272.1.223.23.21	9.4

Table 33: Common (continued)

Object Name	Object OID	New in Release
rcNlsServiceProbeDnsServerListEntry	1.3.6.1.4.1.2272.1.223.26.1	9.4
rcNlsServiceProbeDnsServerListIp	1.3.6.1.4.1.2272.1.223.26.1.2	9.4
rcNlsServiceProbeDnsServerListQueryStatus	1.3.6.1.4.1.2272.1.223.26.1.3	9.4
rcNlsServiceProbeDnsServerListTable	1.3.6.1.4.1.2272.1.223.26	9.4
rcNlsServiceProbeDnsServerListType	1.3.6.1.4.1.2272.1.223.26.1.1	9.4

Table 34: VSP 4900 Series

Object Name	Object OID	New in Release
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
rcAutoSenseSdWanVrfName	1.3.6.1.4.1.2272.1.231.1.1.1.32	9.0.3
rcCombinedIpAdEntAccessType	1.3.6.1.4.1.2272.1.8.33.1.7	9.4
rcIpAdEntAccessType	1.3.6.1.4.1.2272.1.8.2.1.14	9.4
rcIpBgpExtPeerAfAccessType	1.3.6.1.4.1.2272.1.8.101.16.6.1.41	9.4
rcIpBgpPeerGroupAccessType	1.3.6.1.4.1.2272.1.8.101.11.1.42	9.4
rcIpRouteAccessType	1.3.6.1.4.1.2272.1.8.7.1.17	9.4
rcIpfixNniConfState	1.3.6.1.4.1.2272.1.66.1.1.6	9.4
rcIsisPlsbRoutedMulticastTtl1Bridging	1.3.6.1.4.1.2272.1.63.4.1.22	9.4
rcPortIpfixEnable	1.3.6.1.4.1.2272.1.4.10.1.1.144	9.4
rcSysPtpEnable	1.3.6.1.4.1.2272.1.1.131	9.4
rcVrfAccessType	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.17	9.4
rcVrfBgpOrigin	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.21	9.4
rcVrfIpVpnIsidUnrestrictedNumber	1.3.6.1.4.1.2272.1.203.1.1.4.1.11	9.4
rcVrfIpVpnIsidUntrustedNumber	1.3.6.1.4.1.2272.1.203.1.1.4.1.12	9.4
rcVrfOrigin	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.14	9.4
rcVrfSegmented	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.16	9.4

Table 34: VSP 4900 Series (continued)

Object Name	Object OID	New in Release
rcVrfTrustedVrfId	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.18	9.4
rcVrfUnrestrictedVrfId	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.19	9.4
rcVrfUntrustedVrfId	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.20	9.4

Table 35: VSP 7400 Series

Object Name	Object OID	New in Release
rcnRateLimitExceededTrap	1.3.6.1.4.1.2272.1.21.0.369	9.3
rcRateLimitAction	1.3.6.1.4.1.2272.1.14.25	9.3
rcRateLimitActionPollInterval	1.3.6.1.4.1.2272.1.14.25.1	9.3
rcRateLimitActionTable	1.3.6.1.4.1.2272.1.14.25.2	9.3
rcRateLimitActionEntry	1.3.6.1.4.1.2272.1.14.25.2.1	9.3
rcRateLimitActionIfIndex	1.3.6.1.4.1.2272.1.14.25.2.1.1	9.3
rcRateLimitIfActionShutdownTimeout	1.3.6.1.4.1.2272.1.14.25.2.1.2	9.3
rcRateLimitIfActionTrapInterval	1.3.6.1.4.1.2272.1.14.25.2.1.3	9.3
rcCombinedIpAdEntAccessType	1.3.6.1.4.1.2272.1.8.33.1.7	9.4
rcIpAdEntAccessType	1.3.6.1.4.1.2272.1.8.2.1.14	9.4
rcIpBgpExtPeerAfAccessType	1.3.6.1.4.1.2272.1.8.101.16.6.1.41	9.4
rcIpBgpPeerGroupAccessType	1.3.6.1.4.1.2272.1.8.101.11.1.42	9.4
rcIpRouteAccessType	1.3.6.1.4.1.2272.1.8.7.1.17	9.4
rcIpfixNniConfState	1.3.6.1.4.1.2272.1.66.1.1.6	9.4
rcIsisPlsbRoutedMulticastTtl1Bridging	1.3.6.1.4.1.2272.1.63.4.1.22	9.4
rcPortIpfixEnable	1.3.6.1.4.1.2272.1.4.10.1.1.144	9.4
rcSysPtpEnable	1.3.6.1.4.1.2272.1.1.131	9.4
rcVossSystemFanMinSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.12	9.4
rcVrfAccessType	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.17	9.4
rcVrfBgpOrigin	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.21	9.4
rcVrfIpVpnIsidUnrestrictedNumber	1.3.6.1.4.1.2272.1.203.1.1.4.1.11	9.4
rcVrfIpVpnIsidUntrustedNumber	1.3.6.1.4.1.2272.1.203.1.1.4.1.12	9.4
rcVrfOrigin	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.14	9.4
rcVrfSegmented	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.16	9.4
rcVrfTrustedVrfId	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.18	9.4
rcVrfUnrestrictedVrfId	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.19	9.4
rcVrfUntrustedVrfId	1.3.6.1.4.1.2272.1.203.1.1.1.2.1.20	9.4