# Fabric Engine Release Notes

For Fabric Engine Release 8.6

# Table of Contents

# About this Document

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

## Purpose

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

This document includes the following information:

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

## Conventions

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

## Text Conventions

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

**Table 1: Notes and warnings**

| Icon | Notice type | Alerts you to... |
|------|-------------|------------------|
| | Tip | Helpful tips and notices for using the product. |
| | Note | Useful information or instructions. |
| | Important | Important features or instructions. |
| | Caution | Risk of personal injury, system damage, or loss of data. |
| | Warning | Risk of severe personal injury. |

## Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation

Release Notes

Hardware and software compatibility for Extreme Networks products

Extreme Optics Compatibility

Other resources such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

## Help and Support

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

Extreme Portal

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

The Hub

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

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

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

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

## Subscribe to Product Announcements

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

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

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

# Send Feedback

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

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

To send feedback, do either of the following:

- Access the feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.
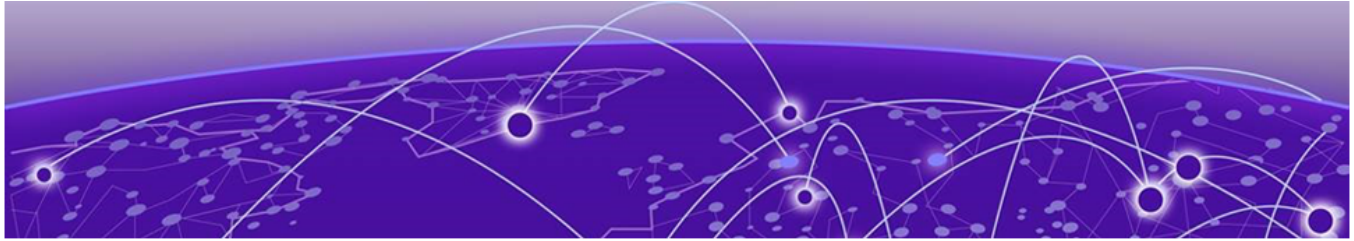
# New in this Release

> **Note**
>
> Beginning with this release, the names of the network operating systems have changed on the universal hardware. VSP Operating System Software (VOSS) has been renamed to Fabric Engine and ExtremeXOS has been renamed to Switch Engine.
>
> For non-universal hardware switches, the names of the network operating systems will continue to be VSP Operating System Software (VOSS) and ExtremeXOS.

The following platforms support Fabric Engine 8.6:

* ExtremeSwitching 5320 Series
* ExtremeSwitching 5420 Series
* ExtremeSwitching 5520 Series

> **Note**
>
> Upgrading the firmware from an earlier version of VOSS to Fabric Engine 8.6, or later, on the ExtremeSwitching 5420 and 5520 series will change the SNMP SysObjectID value. This change might affect SNMP-based management systems. For more information, see this Knowledge Article.

## New Hardware

### 5320 Series

The ExtremeSwitching 5320 Series is a family of feature-rich edge switches. The 5320 Series universal hardware provides end-to-end secure network segmentation, in addition to advanced policy capabilities, and offers a user-selectable choice of Fabric Engine and Switch Engine network operating system (NOS) personalities.

Fabric Engine 8.6 supports the following new switching models:

- 5320-24P-8XE

  ◦ 24 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) ports with 802.3af Type 1 PoE and 802.3at Type 2 PoE+
  ◦ 8 10G SFP+ ports
  ◦ All ports support MACsec.
  ◦ Ports 29-32 support 100M, 1G, and 10G optics.
  ◦ By default, ports 1/25-1/27 are reserved for internal loopback function when the switch operates in SPBM mode.
  ◦ In non SPBM mode, ports 25-32 support 1G and 10G optics.

- 5320-24T-8XE

  ◦ 24 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) ports
  ◦ 8 10G SFP+ ports
  ◦ All ports support MACsec.
  ◦ Ports 29-32 support 100M, 1G, and 10G optics.
  ◦ By default, ports 1/25-1/27 are reserved for internal loopback function when the switch operates in SPBM mode.
  ◦ In non SPBM mode, ports 25-32 support 1G and 10G optics.

- 5320-48P-8XE

  ◦ 48 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) ports with 802.3af Type 1 PoE and 802.3at Type 2 PoE+
  ◦ 8 10G SFP+ ports
  ◦ All ports support MACsec.
  ◦ Ports 53-56 support 100M, 1G, and 10G optics.
  ◦ By default, ports 1/49-1/51 are reserved for internal loopback function when the switch operates in SPBM mode.
  ◦ In non SPBM mode, ports 49-56 support 1G and 10G optics.

- 5320-48T-8XE

  ◦ 48 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) ports
  ◦ 8 10G SFP+ ports
  ◦ All ports support MACsec.
  ◦ Ports 53-56 support 100M, 1G, and 10G optics.
  ◦ By default, ports 1/49-1/51 are reserved for internal loopback function when the switch operates in SPBM mode.
  ◦ In non SPBM mode, ports 49-56 support 1G and 10G optics.

A future release will add support for the following two models, which are included in the current documentation:

- 5320-16P-4XE

  ◦ 16 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) ports with 802.3af Type 1 PoE and 802.3at Type 2 PoE+
  ◦ 4 10G SFP+ ports

- ◦ Powered by AC power source
- ◦ All ports support MACsec.
- 5320-16P-4XE-DC
  - ◦ 16 10/100/1000BASE-T full-duplex (FDX), half-duplex (HDX) ports with 802.3af Type 1 PoE and 802.3at Type 2 PoE+
  - ◦ 4 10G SFP+ ports
  - ◦ Powered by DC power source
  - ◦ All ports support MACsec.

In addition to the fixed ports, all models provide console interface ports (one micro Type B USB and one RJ-45) and one Type A USB port for removable storage.

The 5320 Series family of switches have either 4 or 8 SFP+ ports:

- All SFP+ ports support 1G and 10G speeds.
- The last four SFP+ ports of the 24- and 48-port models also support 100M.

For optics compatibility, see the Extreme Optics website.

For 10G operation, you must install a port license.

For information about feature support, see *Fabric Engine Feature Support Matrix* and *Fabric Engine User Guide*.

## New Transceivers and Components

Extreme Networks might announce support for additional optical components in the future. For a complete and current list of supported optical components, see the Extreme Optics website.

## 100 Mb on 5520-24X

This release introduces 100 Mb support on 5520-24X switches.

# New Software Features or Enhancements

The following sections describe what is new in this release:

## Fabric Extend Enhancement

In earlier releases, you could not configure the tunnel source address (`ip-tunnel-source-address` command) if the address belonged to a VRF with an attached I-SID. You can now use an *overlay* parameter with this command to bypass this restriction.

For more information, see *Fabric Engine User Guide*.

## Fail Open Continuity Mode

Continuity Mode is an enhancement to existing Fail Open VLAN and Fail Open I-SID behavior that shortens the transition to the Fail Open VLAN or Fail Open I-SID. This mode is automatically enabled if

you configure Fail Open VLAN or Fail Open I-SID; you cannot disable Continuity Mode. Continuity Mode replaces the Fail Open behavior of previous releases.

For more information, see *Fabric Engine User Guide*.

## Increase the Number of SPB Nodes per Area

This release introduces a new boot config flag, which is called `boot config flags spbm-node-scaling`. This command is only applicable to, and supported on, 5320 Series and 5420 Series.

For more information, see *Fabric Engine User Guide* .

## RESTCONF APIs Enhancement

In this release, RESTCONF API support is enhanced to include:

- Get all client MAC addresses (mgmt VRF, GRT and user VRF) from the device.
- Get all switch client ARP information.

For more information, see *Fabric Engine User Guide*.

# Other Changes

## EDM Changes

This release makes the following navigational changes in EDM:

| Previous release | Current release |
|---|---|
| **Configuration** > **Edit** > **Diagnostics** | Moved to **Configuration** > **Serviceability** > **Diagnostics** |
| **Configuration** > **IS-IS** | Renamed to **Configuration** > **Fabric** |
| **Configuration** > **Edit** > **Fabric Attach** | Moved to **Configuration** > **Fabric** > **Fabric Attach** |
| **Configuration** > **Edit** > **DVR** | Moved to **Configuration** > **Fabric** > **DVR** |
| **Configuration** > **Edit** > **AutoSense** | Moved to **Configuration** > **Fabric** > **AutoSense** |

# Filenames for this Release

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

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

When extracting the software image file, the extraction process appends the software version portion of the extracted filenames to include the final full software version. (For example, extracting

**5520.8.2.5.0.voss** results in a software file named **5520.8.2.5.0.GA**.) Ensure that you specify the final full software version (in this case, **8.2.5.0.GA**) when using CLI commands that include the software version, such as activating or removing the software.

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

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

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

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

**Table 2: 5320 Series Software Filenames and Sizes**

| Description | File | Size |
| --- | --- | --- |
| SHA512 Checksum files | 5320.8.6.0.0.sha512 | 1,526 bytes |
| MD5 Checksum files | 5320.8.6.0.0.md5 | 512 bytes |
| MIB - supported object names | 5320.8.6.0.0_mib_sup.txt | 1,525,513 bytes |
| MIB - zip file of all MIBs | 5320.8.6.0.0_mib.zip | 1,219,035 bytes |
| MIB - objects in the OID compile order | 5320.8.6.0.0_mib.txt | 8,156,676 bytes |
| EDM Help files | FabricEnginev860_HELP_EDM_gzip.zip | 4,613,738 bytes |
| Logs reference | 5320.8.6.0.0_edoc.tar | 64,020,480 bytes |
| Software image | 5320.8.6.0.0.voss | 102,452,958 bytes |
| Open source software - Master copyright file | 5320.8.6.0.0_oss-notice.html | 2,785,868 bytes |
| YANG model | restconf_yang.tgz | 506,020 bytes |

**Table 3: 5420 Series Software Filenames and Sizes**

| Description | File | Size |
| --- | --- | --- |
| SHA512 Checksum files | 5420.8.6.0.0.sha512 | 1,526 bytes |
| MD5 Checksum files | 5420.8.6.0.0.md5 | 512 bytes |
| MIB - supported object names | 5420.8.6.0.0_mib_sup.txt | 1,524,589 bytes |
| MIB - zip file of all MIBs | 5420.8.6.0.0_mib.zip | 1,219,035 bytes |
| MIB - objects in the OID compile order | 5420.8.6.0.0_mib.txt | 8,156,676 bytes |
| EDM Help files | FabricEnginev860_HELP_EDM_gzip.zip | 4,613,738 bytes |
| Logs reference | 5420.8.6.0.0_edoc.tar | 64,020,480 bytes |
| Software image | 5420.8.6.0.0.voss | 102,042,958 bytes |

Table 3: 5420 Series Software Filenames and Sizes (continued)

| Description | File | Size |
|---|---|---|
| Open source software - Master copyright file | 5420.8.6.0.0_oss-notice.html | 2,785,868 bytes |
| YANG model | restconf_yang.tgz | 506,020 bytes |

Table 4: 5520 Series Software Filenames and Sizes

| Description | File | Size |
|---|---|---|
| SHA512 Checksum files | 5520.8.6.0.0.sha512 | 1,526 bytes |
| MD5 Checksum files | 5520.8.6.0.0.md5 | 512 bytes |
| MIB - supported object names | 5520.8.6.0.0_mib_sup.txt | 1,525,146 bytes |
| MIB - zip file of all MIBs | 5520.8.6.0.0_mib.zip | 1,219,035 bytes |
| MIB - objects in the OID compile order | 5520.8.6.0.0_mib.txt | 8,156,676 bytes |
| EDM Help files | FabricEnginev860_HELP_EDM_gzip.zip | 4,613,738 bytes |
| Logs reference | 5520.8.6.0.0_edoc.tar | 64,020,480 bytes |
| Software image | 5520.8.6.0.0.voss | 106,095,381 bytes |
| Open source software - Master copyright file | 5520.8.6.0.0_oss-notice.html | 2,785,868 bytes |
| YANG model | restconf_yang.tgz | 506,020 bytes |

# Upgrade and Downgrade Considerations

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

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

See the *Fabric Engine User Guide* for detailed image management procedures that includes information about the following specific upgrade considerations:

- IPv6:
  ◦ Notes for systems using IPv6 static neighbors

## Validated Upgrade Paths

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

### Validated Upgrade Path

Validated upgrade path for 5420 Series and 5520 Series:

- VOSS 8.5.x to Fabric Engine 8.6
- VOSS 8.4.x to Fabric Engine 8.6

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

### Switches That Will Not Use Zero Touch Deployment

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

1. Upgrade to this release from one of the previously described releases.

2. Continue to use the previous switch configuration.

## Switches That Will Use Zero Touch Deployment

Switches that will use Zero Touch Deployment with ExtremeCloud IQ or ZTP+ with Extreme Management Center or 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 Extreme Management Center or ExtremeCloud IQ - Site Engine; use the upgrade functionality available in ExtremeCloud IQ or Extreme Management Center or ExtremeCloud IQ - Site Engine.

1. Upgrade to this release from one of the previously described releases.
2. Ensure the switch boots without a configuration file. To ensure the switch boots without a configuration file, perform one of the following actions:

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

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

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

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

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

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

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

3. Reboot the switch.

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

- The ssh and sshd boot configuration flags are enabled by default.
- All ports are Private VLAN isolated ports.
- VLAN 4048 is created as an *onboarding-vlan* for host-only connectivity for In Band management. All front panel ports are members of VLAN 4048.
- In Band management is enabled.
- Dynamic Host Configuration Protocol (DHCP) client requests are cycled between In Band and Out of Band ports, except on the 5320 Series. 5320 Series supports In Band management only.

- If the switch resets after the IP address is obtained from the DHCP Server, the entire DHCP process does not need to be repeated. Instead, the switch can directly send the DHCP Request to the DHCP Server for the IP stored in the /intflash/dhcp/dhclient.leases file.
- Out of Band management is enabled, except on the 5320 Series. 5320 Series supports In Band management only.
- All ports are administratively enabled.
- IQAgent is enabled by default.
- Zero Touch Provisioning Plus (ZTP+) for Extreme Management Center or ExtremeCloud IQ - Site Engine onboarding is enabled by default.
- Initiates Zero Touch Fabric Configuration.
- 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 Extreme Management Center or ExtremeCloud IQ - Site Engine.

All switches also receive a Zero Touch Fabric Configuration. For more information, see *Fabric Engine User Guide*.

## Downgrade Considerations

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

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

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

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

For information about how to reinstall ExtremeCloud IQ Agent firmware, see *Fabric Engine User Guide*.

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

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

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

To add new Zero Touch Fabric Configuration devices or implement Zero Touch Fabric Configuration on existing devices, the network requires a nickname server and reachability to the DHCP server and, optionally, ExtremeCloud IQ servers or Extreme Management Center or ExtremeCloud IQ - Site Engine. How you implement this depends on if the network is a new deployment, or an existing Fabric network that you upgrade. In a new deployment, you can meet the network requirements with one node, known as a seed node. In an existing network, functions can already exist on different nodes.

For more details on Zero Touch Fabric Configuration, see *Fabric Engine User Guide*.

> **Important**
>
> Not all upgrade paths are validated by Extreme Networks for each new software release. To understand the validated upgrade paths, see Validated Upgrade Paths on page 14.

## Network Requirements

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

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

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

  For Zero Touch Fabric Configuration to work when a new switch that runs a Zero Touch Fabric release, connects to a switch on an existing Fabric, upgrade at least the existing Fabric switches to a Zero Touch Fabric release first.
- Some SPB deployments use Ethertype 0x88a8 but many use 0x8100. Zero Touch Fabric Configuration works with existing networks that use either value as long as the existing switches that connect to the new switches run a Zero Touch Fabric release.

## Zero Touch Fabric Configuration Switch

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

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

1. Upgrade to a Zero Touch Fabric release, if the device is not a new deployment already running a Zero Touch Fabric release. As a best practice, upgrade to a Fabric Engine release. For a new deployment of universal hardware, ensure the network operating system (NOS) is Fabric Engine.

2. On upgraded existing switches, ensure the switch boots without a configuration file. The switch joins the network as an end host. To ensure the switch boots without a configuration file, perform one of the following actions:

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

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

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

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

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

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

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

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

   > **Note**
   >
   > For more details on Zero Touch Deployment, see *Fabric Engine User Guide*.

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

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

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

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

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

- Enables Auto-sense on all ports.
- Configures Auto-sense access ports and Layer 2 trusted Auto-sense ports.
- Enables IS-IS globally.
- With Auto-sense, ports on a switch can detect whether they connect to an SPB device, a Fabric Attach (FA) client, FA Proxy, Voice IP devices, or an undefined host, and then make the necessary configuration.

4. If the seed node uses Auto-sense IS-IS Authentication, configure the remote switch to use the same authentication type and key as the seed node.
5. The switch joins the Fabric.
6. The nickname server dynamically assigns an SPBM nickname.
7. After the Zero Touch Fabric establishes successfully, the switch attempts to acquire an IP address on the onboarding VLAN and I-SID using DHCP. When the DHCP client obtains an IP address for the switch, the switch automatically attempts to connect to ExtremeCloud IQ and Extreme Management Center or ExtremeCloud IQ - Site Engine.

# Hardware and Software Compatibility

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

## 5320 Series Hardware

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

**Table 5: Switch models**

| Model part number | Initial release | Supported new feature release |
|---|---|---|
| | | Fabric Engine 8.6 |
| 5320-16P-4XE | Future release | N |
| 5320-16P-4XE-DC | Future release | N |
| 5320-24P-8XE | Fabric Engine 8.6 | Y |
| 5320-24T-8XE | Fabric Engine 8.6 | Y |
| 5320-48P-8XE | Fabric Engine 8.6 | Y |
| 5320-48T-8XE | Fabric Engine 8.6 | Y |

## 5420 Series Hardware

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

**Note**

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

**Table 6: Switch models**

| Model part number | Initial release | Supported new feature release | | |
|---|---|---|---|---|
| | | VOSS 8.4.2 | VOSS 8.5 | Fabric Engine 8.6 |
| 5420F-24T-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-8W-16P-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-24P-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-48T-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-16MW-32P-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-24S-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-16W-32P-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-48P-4XE | VOSS 8.4 | Y | Y | Y |
| 5420F-48P-4XL | VOSS 8.4 | Y | Y | Y |
| 5420M-24T-4YE | VOSS 8.4 | Y | Y | Y |
| 5420M-24W-4YE | VOSS 8.4 | Y | Y | Y |
| 5420M-48T-4YE | VOSS 8.4 | Y | Y | Y |
| 5420M-48W-4YE | VOSS 8.4 | Y | Y | Y |
| 5420M-16MW-32P-4YE | VOSS 8.4 | Y | Y | Y |

# 5520 Series Hardware

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

**Note**

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

**Table 7: Switch models**

| Model part number | Initial release | Supported new feature release | | | | |
|---|---|---|---|---|---|---|
| | | VOSS 8.3 | VOSS 8.4 | VOSS 8.4.2 | VOSS 8.5 | Fabric Engine 8.6 |
| 5520-24T | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-24W | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-48T | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-48W | VOSS 8.2.5 | Y | Y | Y | Y | Y |

**Table 7: Switch models (continued)**

| Model part number | Initial release | Supported new feature release | | | | |
|---|---|---|---|---|---|---|
| | | VOSS 8.3 | VOSS 8.4 | VOSS 8.4.2 | VOSS 8.5 | Fabric Engine 8.6 |
| 5520-12MW-36W | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-24X | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-48SE | VOSS 8.2.5 | Y | Y | Y | Y | Y |

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

**Table 8: Versatile Interface Modules (VIMs)**

| Model part number | Initial release | Supported new feature release | | | | |
|---|---|---|---|---|---|---|
| | | VOSS 8.3 | VOSS 8.4 | VOSS 8.4.2 | VOSS 8.5 | Fabric Engine 8.6 |
| 5520-VIM-4X | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-VIM-4XE | VOSS 8.2.5 | Y | Y | Y | Y | Y |
| 5520-VIM-4YE | VOSS 8.2.5 | Y | Y | Y | Y | Y |

## Versatile Interface Module Operational Notes

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

**Table 9: 5520-VIM Matrix**

| | 5520-VIM-4X | 5520-VIM-4XE | 5520-VIM-4YE |
|---|---|---|---|
| Operational speeds | 1Gbps & 10Gbps | 1Gbps & 10Gbps | 10Gbps & 25Gbps |
| PHY present | No | Yes | Yes |
| 1000BASE-T & 10GBASE-T | 10GBASE-T only | Both | 10GBASE-T only |
| Mixed speeds | 1Gbps & 10Gbps | 1Gbps & 10Gbps | Mixed speeds not supported |
| 1G Auto-negotiation | Disabled | Disabled | Disabled |
| 10G Auto-negotiation | Disabled | Disabled | Disabled |
| 25G Auto-negotiation | | | Enabled for DAC Disabled for Fiber |
| FEC | Not supported | Not supported | Auto-FEC enabled for DAC and Fiber |
| MACsec | Not supported | 128/256 bit | 128/256 bit |

## Operational Notes for VIM Transceivers

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

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

# Transceivers

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

To find product descriptions and compatibility information for optical transceivers and components, visit the Extreme Optics website.

## Auto-Negotiation

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

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

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

## Forward Error Correction (FEC)

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

For more information about FEC, see *Fabric Engine User Guide*.

# Power Supply Compatibility

You can use certain power supplies in more than one platform. This section lists the power supplies and indicates the compatible platforms.

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

- *ExtremeSwitching 5320 Series Hardware Installation Guide*
- *ExtremeSwitching 5420 Series Hardware Installation Guide*
- *ExtremeSwitching 5520 Series Hardware Installation Guide*

Most 5320 Series switches have one fixed AC power supply unit. The 5320-16P-4XE-DC switch has one fixed DC power supply unit.

In the following table for 5420 Series power supplies:

- All 5420-compatible power supplies have front-to-back ventilation airflow.
- Each power supply has a keyed power inlet (C16) that requires a notched (C15) power cord.

**Table 10: 5420 Series Power Supplies**

| Platform | 150W AC PSU-FB XN-ACPWR-150W | 600W AC PoE-FB XN-ACPWR-600W | 920W AC PoE XN-ACPWR-920W | 1200 W AC PoE-FB XN-ACPWR-1200W | 2000 W AC PSU-FB XN-ACPWR-2000W-F |
|---|---|---|---|---|---|
| 5420F-24T-4XE | Y | | | | |
| 5420F-24S-4XE | Y | | | | |
| 5420F-8W-16P-4XE | | Y | | | |
| 5420F-24P-4XE | | Y | | | |
| 5420F-48T-4XE | Y | | | | |
| 5420F-16MW-32P-4XE | | | | Y | |
| 5420F-16W-32P-4XE | | | | Y | |
| 5420F-48P-4XE | | | Y | | |
| 5420F-48P-4XL | | | Y | | |
| 5420M-24T-4YE | Y | | | | |
| 5420M-24W-4YE | | | Y | | |
| 5420M-48T-4YE | Y | | | | |
| 5420M-48W-4YE | | | Y | | Y |
| 5420M-16MW-32P-4YE | | | Y | | Y |

**Note**

In the following table for 5520 Series power supplies:

- All 5520-compatible power supplies have front-to-back ventilation airflow.
- Each power supply has a keyed power inlet (C16) that requires a notched (C15) power cord.

**Table 11: 5520 Series Power Supplies**

| Platform | 350 W AC PSU-FB 10953 | 715 W AC PSU-FB 10951 | 1100 W AC PSU-FB 10941 | 2000 W AC PSU-FB XN-ACPWR-2000W-F |
|---|---|---|---|---|
| 5520-24T | Y | | | |
| 5520-24W | | Y | Y | Y |
| 5520-24X | Y | | | |
| 5520-12MW-36W | | Y | Y | Y |
| 5520-48T | Y | | | |
| 5520-48W | | Y | Y | Y |
| 5520-48SE | Y | | | |
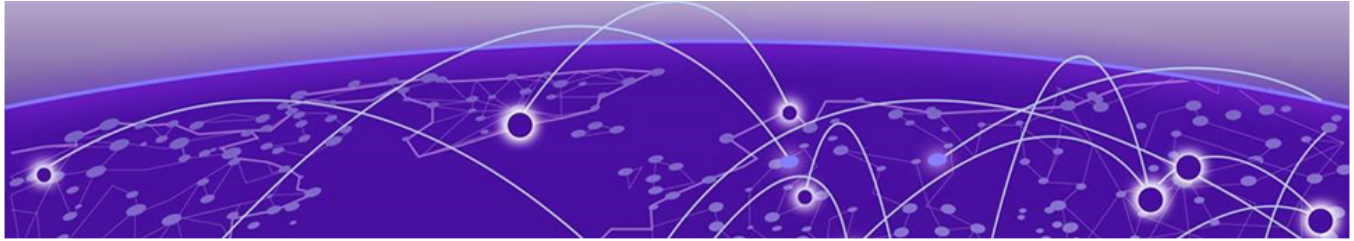
# Scaling

This section documents scaling capabilities of the universal hardware platforms.

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

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

# Layer 2

Table 12: Layer 2 Maximums

| Attribute | Product | Maximum number supported |
|---|---|---|
| MAC table size (without SPBM) | 5320 Series | 32,000 |
| | 5420 Series | 32,000 for 5420F Series models<br>64,000 for 5420M Series models |
| | 5520 Series | 81,920 |
| MAC table size (with SPBM) | 5320 Series | 16,000 |
| | 5420 Series | 16,000 for 5420F Series models<br>32,000 for 5420M Series models |
| | 5520 Series | 40,960 |
| Endpoint Tracking MAC addresses per switch | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 8,000 |
| Directed Broadcast interfaces | 5320 Series | 200<br>See Maximum Number of Directed Broadcast Interfaces on page 29. |
| | 5420 Series | 200<br>See Maximum Number of Directed Broadcast Interfaces on page 29. |
| | 5520 Series | 200<br>See Maximum Number of Directed Broadcast Interfaces on page 29. |
| Port-based VLANs<br><br>**Note:**<br>When you use Flex-UNI functionality, you can use the complete range from 1 to 4096 for port VLAN IDs. | 5320 Series | 4,059 |
| | 5420 Series | 4,059 |
| | 5520 Series | 4,059 |
| Private VLANs | 5320 Series | 50 |
| | 5420 Series | 100 |
| | 5520 Series | 200 |
| Protocol-based VLANs (IPv6 only) | 5320 Series | 1 |
| | 5420 Series | 1 |
| | 5520 Series | 1 |

Table 12: Layer 2 Maximums (continued)

| Attribute | Product | Maximum number supported |
|---|---|---|
| RSTP instances | 5320 Series | 1 |
| | 5420 Series | 1 |
| | 5520 Series | 1 |
| MSTP instances | 5320 Series | 12 |
| | 5420 Series | 12 |
| | 5520 Series | 12 |
| LACP aggregators | 5320 Series | 48-port models: 56 24-port models: 32 16-port models: 20 |
| | 5420 Series | 56 (48 fixed ports, 4 Universal Ethernet ports, 4 SFP-DD ports) |
| | 5520 Series | 48-port models: up to 60 with channelization 24-port models: up to 36 with channelization |
| Ports per LACP aggregator | 5320 Series | 8 active |
| | 5420 Series | 8 active |
| | 5520 Series | 8 active |
| MLT groups | 5320 Series | 48-port models: 56 24-port models: 32 16-port models: 20 |
| | 5420 Series | 56 (48 fixed ports, 4 Universal Ethernet ports, 4 SFP-DD ports) |
| | 5520 Series | 48-port models: up to 60 with channelization 24-port models: up to 36 with channelization |
| Ports per MLT group | 5320 Series | 8 active |
| | 5420 Series | 8 active |
| | 5520 Series | 8 |
| Link State Tracking (LST) groups | 5320 Series | 48 |
| | 5420 Series | 48 |
| | 5520 Series | 48 |

**Table 12: Layer 2 Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| Interfaces per LST group | 5320 Series | 48-port models: 9 upstream/128 downstream<br>16- and 24-port models: 8 upstream/128 downstream |
| | 5420 Series | 8 upstream<br>128 downstream |
| | 5520 Series | 8 upstream<br>128 downstream |
| SLPP VLANs | 5320 Series | 128 |
| | 5420 Series | 128 |
| | 5520 Series | 128 |
| VLACP interfaces | 5320 Series | 48-port models: 56<br>24-port models: 32<br>16-port models: 20 |
| | 5420 Series | 56 (48 fixed ports, 4 Universal Ethernet ports, 4 SFP-DD ports) |
| | 5520 Series | 48-port models: up to 60 with channelization<br>24-port models: up to 36 with channelization |
| Microsoft NLB cluster IP interfaces | 5320 Series | Not supported |
| | 5420 Series | Not supported |
| | 5520 Series | 200<br>See Maximum Number of Microsoft NLB Cluster IP Interfaces on page 29. |

## Maximum Number of Directed Broadcast Interfaces

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

## Maximum Number of Microsoft NLB Cluster IP Interfaces

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

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

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

# IP Unicast

**Table 13: IP Unicast Maximums**

| Attribute | Product | Maximum number supported |
|---|---|---|
| IP interfaces (IPv4 or IPv6 or IPv4+IPv6) | 5320 Series | 248<br>See IP Interface Maximums for 5320 Series on page 35. |
| | 5420 Series | 248<br>See IP Interface Maximums for 5420 Series on page 35. |
| | 5520 Series | 500<br>See IP Interface Maximums for 5520 Series on page 35. |
| VRRP interfaces (IPv4 or IPv6) | 5320 Series | 48-port models: 124<br>16- and 24-port models: 64<br>See IP Interface Maximums for 5320 Series on page 35. |
| | 5420 Series | 124<br>See IP Interface Maximums for 5420 Series on page 35. |
| | 5520 Series | 252<br>See IP Interface Maximums for 5520 Series on page 35. |
| Routed Split Multi-Link Trunking (RSMLT) interfaces (IPv4 or IPv6 or IPv4+IPv6) | 5320 Series | n/a |
| | 5420 Series | 124<br>See IP Interface Maximums for 5420 Series on page 35. |
| | 5520 Series | 499<br>See IP Interface Maximums for 5520 Series on page 35. |

**Table 13: IP Unicast Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| VRRP interfaces with fast timers (200ms) - IPv4/IPv6 | 5320 Series | 24 |
| | 5420 Series | 24 |
| | 5520 Series | 24 |
| ECMP groups/paths per group | 5320 Series | 48-port models: 64/8<br>16- and 24-port models:32 /8 |
| | 5420 Series | 64/8 |
| | 5520 Series | 125/8 |
| OSPF v2/v3 interfaces | 5320 Series | 48-port models: 50<br>16- and 24-port models: 1 |
| | 5420 Series | 50 |
| | 5520 Series | 100 |
| OSPF v2/v3 neighbors (adjacencies) | 5320 Series | 50 |
| | 5420 Series | 50 |
| | 5520 Series | 100 |
| OSPF areas | 5320 Series | 48-port models: 12<br>16- and 24-port models: 4 |
| | 5420 Series | 12 for the switch |
| | 5520 Series | 12 for each VRF<br>80 for the switch |
| IPv4 ARP table | 5320 Series | 48-port models: 15,000<br>16- and 24-port models: 8,000 |
| | 5420 Series | 15,000 for 5420F Series models<br>24,000 for 5420M Series models |
| | 5520 Series | 16,000 |
| IPv4 CLIP interfaces | 5320 Series | 64 |
| | 5420 Series | 64 |
| | 5520 Series | 64 |
| IPv4 RIP interfaces | 5320 Series | 50 |
| | 5420 Series | 50 |
| | 5520 Series | 100 |
| IPv4 BGP peers | 5320 Series | 8 |
| | 5420 Series | 8 |
| | 5520 Series | 16 |

Table 13: IP Unicast Maximums (continued)

| Attribute | Product | Maximum number supported |
|---|---|---|
| IPv4 VRFs with iBGP | 5320 Series | 48-port models: 8<br>16- and 24-port models: 1 |
| | 5420 Series | 8 |
| | 5520 Series | 16 |
| IPv4/IPv6 VRF instances<br>For additional information, see VRF Scaling on page 53. | 5320 Series | 48-port models: 64<br>16- and 24-port models: 1<br>See IP Interface Maximums for 5320 Series on page 35. |
| | 5420 Series | 64<br>See IP Interface Maximums for 5420 Series on page 35. |
| | 5520 Series | 256 including mgmt VRF and GRT<br>See IP Interface Maximums for 5520 Series on page 35. |
| IPv4 static ARP entries | 5320 Series | 48-port models: 1,000 per VRF/5,000 per switch<br>16- and 24-port models: 1,000 per switch |
| | 5420 Series | 1,000 per VRF<br>5,000 per switch |
| | 5520 Series | 2,000 for each VRF<br>10,000 for the switch |
| IPv4 static routes | 5320 Series | 48-port models: 500 per VRF/2,500 per switch<br>16- and 24-port models: 500 per switch |
| | 5420 Series | 500 per VRF<br>2500 per switch |
| | 5520 Series | 1,000 for each VRF<br>5,000 for the switch |
| IPv4 route policies | 5320 Series | 48-port models: 50 per VRF/500 per switch<br>16- and 24-port models: 500 per switch |
| | 5420 Series | 50 per VRF<br>500 per switch |
| | 5520 Series | 500 for each VRF<br>5,000 for the switch |

**Table 13: IP Unicast Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| IPv4 UDP forwarding entries | 5320 Series | 128 |
| | 5420 Series | 128 |
| | 5520 Series | 256 |
| IPv4 DHCP Relay forwarding entries | 5320 Series | 248 |
| | 5420 Series | 248 |
| | 5520 Series | 512 |
| IPv6 DHCP Snoop entries in Source Binding Table | 5320 Series | 48-port models: 513<br>16- and 24-port models: 512 |
| | 5420 Series | 512 |
| | 5520 Series | 1,024 |
| IPv6 Neighbor table | 5320 Series | 8,000 |
| | 5420 Series | 8,000 for 5420F Series models<br>16,000 for 5420M Series models |
| | 5520 Series | 16,000 |
| IPv6 static entries in Source Binding Table | 5320 Series | 48-port models: 65 per VRF/ 256 per switch<br>16- and 24-port models: 256 per switch |
| | 5420 Series | 64 per VRF<br>256 per system |
| | 5520 Series | 128 per VRF<br>512 per system |
| IPv6 static neighbor records | 5320 Series | 48-port models: 64 per VRF/256 per switch<br>16- and 24-port models: 256 per switch |
| | 5420 Series | 64 per VRF<br>256 per switch |
| | 5520 Series | 128 per VRF<br>512 per system |
| IPv6 CLIP interfaces | 5320 Series | 64 |
| | 5420 Series | 64 |
| | 5520 Series | 64 |

Table 13: IP Unicast Maximums (continued)

| Attribute | Product | Maximum number supported |
|---|---|---|
| IPv6 static routes | 5320 Series | 48-port models: 501 16- and 24-port models: 500 |
| | 5420 Series | 500 |
| | 5520 Series | 1,000 |
| IPv6 6in4 configured tunnels | 5320 Series | 32 |
| | 5420 Series | 32 |
| | 5520 Series | 64 |
| IPv6 DHCP Relay forwarding | 5320 Series | 248 |
| | 5420 Series | 248 |
| | 5520 Series | 256 per switch 10 per VRF |
| IPv6 BGP peers | 5320 Series | 8 |
| | 5420 Series | 8 |
| | 5520 Series | 16 Up to 8,000 IPv6 prefixes for BGPv6 peering |
| IPv6 VRFs with iBGP | 5320 Series | 48-port models: 8 16- and 24-port models: 1 |
| | 5420 Series | 8 |
| | 5520 Series | 16 |
| BFD VRF instances | 5320 Series | 48-port models: 16 16- and 24- port models: 1 |
| | 5420 Series | 16 |
| | 5520 Series | 16 |
| BFD sessions per switch (IPv4/IPv6) with default values | 5320 Series | 48-port models: 16 16- and 24- port models: 1 |
| | 5420 Series | 16 |
| | 5520 Series | 16 |
| BFD sessions with Fabric Extend tunnels (IPv4) | 5320 Series | 48-port models: 16 16- and 24- port models: 1 |
| | 5420 Series | 16 |
| | 5520 Series | 16 |

## IP Interface Maximums for 5520 Series

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

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

## IP Interface Maximums for 5420 Series

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

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

## IP Interface Maximums for 5320 Series

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

*16- and 24-port models*

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

*48-port models*

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

## Layer 3 Route Table Size

Table 14: Layer 3 Route Table Size Maximums

| Attribute | Maximum number supported |
|---|---|
| IPv4 RIP routes | See Route Scaling on page 36. |
| IPv4 OSPF routes | |
| IPv4 BGP routes | |

Table 14: Layer 3 Route Table Size Maximums (continued)

| Attribute | Maximum number supported |
|---|---|
| IPv4 SPB shortcut routes | |
| IPv4 SPB Layer 3 VSN routes | |
| IPv6 OSPFv3 routes - GRT only | |
| IPv6 SPB shortcut routes - GRT only | |
| IPv6 RIPng routes | |

## Route Scaling

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

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

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

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

Table 15: 5320 Series

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

**Note:**

The total number of routes include local routes.

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

Table 16: 5420 Series

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

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

Table 17: 5520 Series

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

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

# IP Multicast

Table 18: IP Multicast Maximums

| Attribute | Product | Maximum number supported |
| --- | --- | --- |
| IGMP/MLD interfaces (IPv4/IPv6) | 5320 Series | 4,000/2,000 |
| | 5420 Series | 4,000/2,000 |
| | 5520 Series | 4,059 |

**Table 18: IP Multicast Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| PIM interfaces (IPv4/IPv6) | 5320 Series | 16 active |
| | 5420 Series | 16 active |
| | 5520 Series | 128 active |
| PIM Neighbors (IPv4/IPv6)  (GRT Only) | 5320 Series | 16 |
| | 5420 Series | 16 |
| | 5520 Series | 128 |
| PIM-SSM static channels (IPv4/IPv6) | 5320 Series | 512 |
| | 5420 Series | 512 |
| | 5520 Series | 4,000 |
| Multicast receivers/IGMP joins (IPv4/IPv6) (per switch) | 5320 Series | 6,000 |
| | 5420 Series | 6,000 |
| | 5520 Series | 6,000 |
| Total multicast routes (S,G,V) (IPv4/IPv6) (per switch) | 5320 Series | 48-port models: 4,000<br>16- and 24-port models: 2,000 |
| | 5420 Series | 4,000 |
| | 5520 Series | 4,000 |
| Total multicast routes (S,G,V) (IPv4) on an SPB-PIM Gateway configured switch | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 4,000 |
| Static multicast routes (S,G,V) (IPv4/IPv6) | 5320 Series | 48-port models: 4,000<br>16- and 24-port models: 2,000 |
| | 5420 Series | 4,000 |
| | 5520 Series | 4,000 |
| Multicast enabled Layer 2 VSN (IPv4) | 5320 Series | 48-port models: 500<br>16- and 24-port models: 250 |
| | 5420 Series | 500 |
| | 5520 Series | 2,000 |

**Table 18: IP Multicast Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| Multicast enabled Layer 3 VSN (IPv4) | 5320 Series | 48-port models: 64 16- and 24-port models: 1 |
| | 5420 Series | 64 |
| | 5520 Series | 256 including mgmt VRF and GRT |
| SPB-PIM Gateway controller S,Gs (source announcements) with MSDP (IPv4) | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 6,000 |
| SPB-PIM Gateway controllers per SPB fabric (IPv4) | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 5 |
| SPB-PIM Gateway nodes per SPB fabric (IPv4) | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 64 |
| SPB-PIM Gateway interfaces per BEB (IPv4) | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 64 |
| PIM neighbors per SPB-PIM Gateway node (IPv4) | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 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 30.

**Table 19: DvR Maximums**

| Attribute | Product | Maximum number supported |
|---|---|---|
| **Note:**<br>• On the DvR leaf, you must enable the VRF scaling boot configuration flag if more than 24 VRFs are required in the DvR domain.<br>• Scaling of a VSP 4450 Series switch controls the scaling of the DvR domain it is in. For example, if a VSP 4450 Series switch is in a DvR domain with other platforms such as 5520 Series and 5420 Series, the scaling of the entire domain is limited to the scaling of the VSP 4450 Series. | | |
| DvR Virtual IP interfaces | 5320 Series | 48-port models: 248<br>16- and 24-port models: n/a |
| | 5420 Series | 247 with VIST<br>248 without VIST |
| | 5520 Series | 499 with vIST<br>500 without vIST |
| DvR domains per SPB fabric | 5320 Series | 16 |
| | 5420 Series | 16 |
| | 5520 Series | 16 |
| Controller nodes per DvR domain with default route inject flag enabled<br>Total number of Controllers per domain cannot exceed 8.<br><br>**Note:**<br>A DvR domain containing only Controller nodes and no Leaf nodes can have more than 8 Controllers per domain. | 5320 Series | n/a |
| | 5420 Series | n/a |
| | 5520 Series | 8 |
| Leaf nodes per DvR domain | 5320 Series | 250 |
| | 5420 Series | 250 |
| | 5520 Series | 250 |
| DvR enabled Layer 2 VSNs | 5320 Series | 48-port models: 248<br>16- and 24-port models: n/a |
| | 5420 Series | 247 with vIST<br>248 without vIST |
| | 5520 Series | 499 with vIST<br>500 without vIST |

**Table 19: 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)<br>If DvR Layer 2 VSNs span DvR domains, and all DvR Controllers have an IP interface on the Layer 2 VSNs, then the DvR host scaling is network-wide, as DvR Controllers will consume as many host routes as there are hosts across all DvR domains. | 5320 Series | 48-port models: 16,000<br>16- and 24-port models: n/a |
| | 5420 Series | 16,000 for 5420F Series models<br>32,000 for 5420M Series models |
| | 5520 Series | 48,000 |

# Filters, QoS, and Security

**Table 20: Filters, QoS, and Security Maximums**

| Attribute | Product | Maximum number supported |
|---|---|---|
| For more information, see Filter Scaling on page 42. | | |
| Total IPv4 Ingress rules/ACEs (Port/VLAN/InVSN based, Security/QoS filters) | 5320 Series | 1,024 |
| | 5420 Series | 1,024 |
| | 5520 Series | 1,024 (512 security and 512 QoS) |
| Total IPv4 Egress rules/ACEs (Port based, Security filters) | 5320 Series | 48-port models: 400, 144 if you enable the **ipv6-egress-filter** or **macsec** boot configuration flag<br>16- and 24-port models: 190, 62 if you enable the **ipv6-egress-filter** or **macsec** boot configuration flag |
| | 5420 Series | 400<br>144 if you enable the **ipv6-egress-filter** or **macsec** boot configuration flag |
| | 5520 Series | 336<br>80 if you enable the **ipv6-egress-filter** boot configuration flag |
| Total IPv6 Ingress rules/ACEs (Port/VLAN/InVSN based, Security filters) | 5320 Series | 1,024 |
| | 5420 Series | 512 |
| | 5520 Series | 512 |

**Table 20: Filters, QoS, and Security Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| Total IPv6 egress rules/ACEs (Port based, Security filters) | 5320 Series | 48-port models: 256, 0 with MACsec<br>16- and 24-port models: 128 , 0 with MACsec |
| | 5420 Series | 256, 0 with MACsec |
| | 5520 Series | 256 |
| EAP and NEAP (clients per port) | 5320 Series | 32/800 NEAP<br>700 with MACsec<br>400 with **spbm-node-scaling** |
| | 5420 Series | 32/800 NEAP<br>700 with MACsec<br>400 with **spbm-node-scaling** |
| | 5520 Series | 32 for EAP<br>4,900 for NEAP |

# Filter Scaling

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

*5520 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. This maximum also implies a port member count of 1 for the inPort ACL. The number of rules consumed by inPort ACL is multiplied by the number of ports to which this ACL applies.

- 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

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

- 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.

- 1024 ingress ACEs:

Theoretical maximum of 1024 implies 1 ingress ACL with 512 security ACEs and 512 QoS ACEs

- Ingress ACEs supported: (512 (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.

*5420 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. This maximum also implies a port member count of 1 for the inPort ACL. The number of rules consumed by inPort ACL is multiplied by the number of ports to which this ACL applies.

- 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

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

- 1024 ingress ACEs:

Theoretical maximum of 1024 implies 1 ingress ACL with 512 security ACEs and 512 QoS ACEs

- Ingress ACEs supported: (512 (security) - # of ACLs) + (512 (QoS) - # of ACLs).

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

- 400 egress ACEs:

Theoretical maximum of 400 implies 1 egress ACL with 400 security ACEs

- Egress ACEs supported: 400 - # of ACLs.

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

*5320 Series*

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
  - 512 ACLs with 1 ACE each that can hold either Security/QoS/both action types or
  - a combination based on the following rule: ( (num ACLs + num ACEs) <= 1024)

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

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

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

- 512 IPv6 ingress ACLs (inPort):
  - 512 ACLs with 1 ACE each that can hold either Security/QoS/both action types or
  - a combination based on the following rule: (num ACLs + num ACEs) <= 1024

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

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

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

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

- 400 egress ACEs

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

*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 21: Routed Private VLANs/E-TREEs Maximums**

| | Private VLAN trunk ports | Routed PVLANs/E-TREEs | IPv4 Egress ACE rules available (No IPv6 egress filter bootflag enabled) | IPv4 Egress ACE rules available (With IPv6 egress filter bootflag enabled) |
|---|---|---|---|---|
| 5320-48T-8XE 5320-48P-8XE | 4 | 10 | 349 | 93 |
| 5320-16P-4XE 5320-16P-4XE-DC 5320-24P-8XE 5320-24T-8XE | 4 | 10 | 139 | 11 |

**Table 21: Routed Private VLANs/E-TREEs Maximums (continued)**

| | 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) |
|---|---|---|---|---|
| 5420 Series | 4 | 10 | 349 | 93 |
| 5520 Series | 4 | 10 | 285 | 29 |

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

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

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

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


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

# OAM and Diagnostics

**Table 22: OAM and Diagnostics Maximums**

| Attribute | Product | Maximum number supported |
|---|---|---|
| EDM sessions | 5320 Series | 5 |
| | 5420 Series | 5 |
| | 5520 Series | 5 |
| FTP sessions (IPv4/IPv6) | 5320 Series | 8 total (4 for IPv4 and 4 for IPv6) |
| | 5420 Series | 8 total (4 for IPv4 and 4 for IPv6) |
| | 5520 Series | 8 total (4 for IPv4 and 4 for IPv6) |
| SSH sessions (IPv4/IPv6) | 5320 Series | 8 total (any combination of IPv4 and IPv6) |
| | 5420 Series | 8 total (any combination of IPv4 and IPv6) |
| | 5520 Series | 8 total (any combination of IPv4 and IPv6) |
| Telnet sessions (IPv4/IPv6) | 5320 Series | 16 total (8 for IPv4 and 8 for IPv6) |
| | 5420 Series | 16 total (8 for IPv4 and 8 for IPv6) |
| | 5520 Series | 16 total (8 for IPv4 and 8 for IPv6) |
| TFTP sessions (IPv4/IPv6) | 5320 Series | 2 total (any combination of IPv4 and IPv6) |
| | 5420 Series | 2 total (any combination of IPv4 and IPv6) |
| | 5520 Series | 2 total (any combination of IPv4 and IPv6) |
| Mirrored ports (source) | 5320 Series | 48-port models: 56<br>24-port models: 32<br>16-port models: 20 |
| | 5420 Series | 56 |
| | 5520 Series | 48-port models: 47 (up to 58 with channelization)<br>24-port models: 23 (up to 34 with channelization) |
| Mirroring ports (destination) | 5320 Series | 4 |
| | 5420 Series | 4 |
| | 5520 Series | 4 |

**Table 22: OAM and Diagnostics Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| Fabric RSPAN Port mirror instances per switch (Ingress only) | 5320 Series | Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror. |
| | 5420 Series | Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror. |
| | 5520 Series | Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror. |
| Fabric RSPAN Flow mirror instances per switch (Ingress only) | 5320 Series | Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets. |
| | 5420 Series | Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets. |
| | 5520 Series | Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets. |
| Fabric RSPAN Monitoring I-SIDs (network value) | 5320 Series | 48-port models: 500 Monitoring I-SIDs across SPB network<br>16 and 24-port models: 250 Monitoring I-SIDs across SPB network |
| | 5420 Series | 500 Monitoring I-SIDs across SPB network |
| | 5520 Series | 1,000 Monitoring I-SIDs across SPB network |
| sFlow sampling limit | 5320 Series | 3,100 samples per second |
| | 5420 Series | 3,100 samples per second |
| | 5520 Series | 3,100 samples per second |
| IPFIX flows | 5320 Series | 48-port models: 9,000<br>16- and 24-port models: n/a |
| | 5420 Series | 9,000 |
| | 5520 Series | 36,863 |

**Table 22: OAM and Diagnostics Maximums (continued)**

| Attribute | Product | Maximum number supported |
|---|---|---|
| Application Telemetry host monitoring - maximum number of monitored hosts<br><br>**Note:**<br>These resources are shared with the IPv4 Filter Ingress rules/ACEs. | 5320 Series | 382 hosts |
| | 5420 Series | 382 hosts |
| | 5520 Series | 382 hosts |

# Fabric Scaling

This section lists the fabric scaling information.

**Table 23: Fabric Maximums**

| Attribute | Product | Maximum number supported (with and without vIST) |
|---|---|---|
| Number of SPB IS-IS areas | 5320 Series | 1 |
| | 5420 Series | 1 |
| | 5520 Series | 1 |
| Number of B-VIDs | 5320 Series | 2 |
| | 5420 Series | 2 |
| | 5520 Series | 2 |
| Maximum number of Physical and Logical (Fabric Extend) NNI interfaces/adjacencies | 5320 Series | 64 |
| | 5420 Series | 50 |
| | 5520 Series | 128 |
| SPBM enabled nodes per area (BEB + BCB) | 5320 Series | 350 with no **spbm-node-scaling**<br>500 with **spbm-node-scaling** |
| | 5420 Series | 350 with no **spbm-node-scaling**<br>500 with **spbm-node-scaling** |
| | 5520 Series | 800 |

Table 23: Fabric Maximums (continued)

| Attribute | Product | Maximum number supported (with and without vIST) |
|---|---|---|
| Number of BEBs not part of vIST clusters this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI) | 5320 Series | 350 with no **spbm-node-scaling** 500 with **spbm-node-scaling** |
| | 5420 Series | 350 with no **spbm-node-scaling** 500 with **spbm-node-scaling** |
| | 5520 Series | 800 |
| Number of BEBs that are part of a vIST cluster this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI) | 5320 Series | 300 |
| | 5420 Series | 300 |
| | 5520 Series | 800 |
| I-SIDs supported (local UNI present on device) | 5320 Series | See Number of I-SIDs supported |
| | 5420 Series | See Number of I-SIDs supported |
| | 5520 Series | See Number of I-SIDs supported |
| Maximum number of Layer 2 VSNs per switch (local UNI present on device) | 5320 Series | 48-port models: 500 16- and 24-port models: 250 |
| | 5420 Series | 500 |
| | 5520 Series | 3,580 |
| Maximum number of Switched UNI I-SIDs per switch | 5320 Series | See Number of I-SIDs supported |
| | 5420 Series | See Number of I-SIDs supported |
| | 5520 Series | See Number of I-SIDs supported |
| Maximum number of Transparent Port UNIs per switch | 5320 Series | 48-port models: 53 24-port models: 29 16- models: 20 |
| | 5420 Series | 56 |
| | 5520 Series | 48-port models: 48 24-port models: 24 |
| Maximum number of E-Tree PVLAN UNIs per switch | 5320 Series | 50 |
| | 5420 Series | 100 |
| | 5520 Series | 200 |

**Table 23: Fabric Maximums (continued)**

| Attribute | Product | Maximum number supported (with and without vIST) |
|---|---|---|
| Maximum number of Layer 3 VSNs per switch See VRF Scaling on page 53. | 5320 Series | 48-port models: 64 16- and 24-port models: 1 local VRF and 23 remote accepted I-SIDs |
| | 5420 Series | 64 |
| | 5520 Series | 256 including mgmt VRF and GRT |
| Maximum number of SPB Layer 2 multicast UNI I-SIDs | 5320 Series | See Number of I-SIDs supported |
| | 5420 Series | 500 originated 800 terminated |
| | 5520 Series | See Number of I-SIDs supported |
| Maximum number of SPB Layer 3 multicast UNI I-SIDs | 5320 Series | 500 originated 800 terminated without **spbm-node-scaling** 1,200 with **spbm-node-scaling** |
| | 5420 Series | 500 originated 800 terminated without **spbm-node-scaling** 1,200 with **spbm-node-scaling** |
| | 5520 Series | Maximum 4,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached. |
| Maximum number of FA ISID/VLAN assignments per port | 5320 Series | 94 |
| | 5420 Series | 94 |
| | 5520 Series | 94 |
| Maximum number of IP multicast S,Gs when operating as a BCB (intra-area) | 5320 Series | 16,000 |
| | 5420 Series | 16,000 |
| | | 16,000 |

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

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

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

| Number of IS-IS interfaces (NNIs) | Product | I-SIDs with vIST configured on the platform | I-SIDs without vIST configured on the platform |
|---|---|---|---|
| 4 | 5320 Series | n/a | 500 |
| | 5420 Series | 564 | 564 |
| | 5520 Series | 4,000 | 4,000 |
| 6 | 5320 Series | n/a | 500 |
| | 5420 Series | 564 | 564 |
| | 5520 Series | 3,500 | 4,000 |
| 10 | 5320 Series | n/a | 500 |
| | 5420 Series | 564 | 564 |
| | 5520 Series | 2,900 | 4,000 |
| 20 | 5320 Series | n/a | 500 |
| | 5420 Series | 564 | 564 |
| | 5520 Series | 2,000 | 4,000 |
| 48 | 5320 Series | n/a | 500 |
| | 5420 Series | 564 | 564 |
| | 5520 Series | 1,000 | 2,000 |
| 72 | 5320 Series | n/a | n/a |
| | 5420 Series | n/a | n/a |
| | 5520 Series | 750 | 1,500 |
| 100 | 5320 Series | n/a | n/a |
| | 5420 Series | n/a | n/a |
| | 5520 Series | 550 | 1,100 |
| 128 | 5320 Series | n/a | n/a |
| | 5420 Series | n/a | n/a |
| | 5520 Series | 450 | 900 |
| 250 | 5320 Series | n/a | n/a |
| | 5420 Series | n/a | n/a |
| | 5520 Series | n/a | n/a |

## Interoperability Considerations for IS-IS External Metric

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

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

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

## Recommendations

This section provides recommendations that affect feature configuration.

Pay special attention to the expected scaling of routes in the network and the number of OSPF neighbors in a single VRF when you select configuration values for the **isis 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, the system reserves VLAN IDs 4060 to 4094 for internal use.

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

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

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

# Important Notices

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

## ExtremeCloud IQ Support

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

ExtremeCloud IQ supports the following platforms:

- 5320 Series
- 5420 Series
- 5520 Series

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

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

The switch software integrates with ExtremeCloud IQ using IQAgent.

For more information, see *Fabric Engine User Guide*.

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

## Compatibility with Extreme Management Center or ExtremeCloud IQ - Site Engine

Fabric Engine 8.6 is compatible with ExtremeCloud IQ - Site Engine version 22.3 or later. Older versions
(including Extreme Management Center) will not recognize devices running Fabric Engine.

## Using Ping or IP Traceroute for Hosts in the DvR-One-IP Subnet

To use DvR-One-IP, a circuitless IP (CLIP) must exist in the VRF to which the DvR-One-IP interface
belongs. If the DvR-One-IP interface is part of the global router (GRT), a CLIP must exist in the GRT and
it must be configured as the IS-IS **ip-source-address**. If these CLIPs exist, pinging hosts in the
DvR-One-IP subnet from the DvR Controller works as expected.

If the CLIPs do not exist, pinging hosts in the DvR-One-IP subnet is not possible from DvR Controllers.
The ping attempt times out and the switch displays the following warning message: `Warning: For
DVR one IP a loopback IP must be configured on the VRF`. If you provide a source IP
address with the **ping** command, the switch does not display the warning message but the ping
attempt fails.

This same restriction also applies to IP traceroute.

## AES-GCM SSH Connection with Open SSH

Switch side encryption and authentication type must be set to the AES-GCM-128/256 methods and
needs at least one hmac method in the authentication list in addition for the connection to work.

## Auto Negotiation Settings

The software requires the same auto negotiation settings on link partners to avoid incorrect declaration
of link status. Mismatched settings can cause the links to stay down as well as unpredictable behavior.
Ensure the auto negotiation settings between local ports and their remote link partners match before
upgrading.

## clock set

Configuring time using the **clock set** command on a switch that also uses NTP can cause ports from
going up.

As a best practice, do not configure time using the **clock set** command. Instead, rely on NTP to
synchronize the internal clocks.

## dos-chkdsk

If at the end of the **dos-chkdsk WORD<1-99>** command output you see the following choice:

```
1) Correct

2) Don't correct
```

Then, you should run the **dos-chkdsk WORD<1-99> repair** command.

# Base MAC Address Assignment for Universal Hardware Switches

When running Switch Engine, each universal hardware switch uses a base MAC address at offset 0 for both the default management port, if available, and in-band VLAN utilizing DHCP, for example, 00:c0:cc:8b:68:00. When the switch runs Fabric Engine, it uses a base MAC at offset 0x81 for the default management port (for example, 00:c0:cc:8b:68:81) and offset 256 for the in-band VLAN (for example, 00:c0:cc:8b:69:00).

> **Note**
> The address assignment for the in-band VLAN assumes that the VLAN has a mac-offset value of 0 assigned. If a different mac-offset value is assigned, the MAC address changes accordingly. For example, if mac-offset is 10, then the associated MAC address is 00:c0:cc:8b:69:0A.

When using a DHCP client on the switch, the switch sends a common DHCP client identifier equal to the base MAC address of the switch that is printed on the switch label. Because of this, assuming a standard DHCP pool configuration, the DHCP server always recognizes the switch by the same IP address, regardless of whether Switch Engine or Fabric Engine runs on the switch.

To statically assign IP addresses on the DHCP server, assign them based upon the DHCP client ID. This assignment will ensure that the bindings do not change when the switch alternates between Switch Engine and Fabric Engine. If you assign the DHCP IP addresses based on MAC addresses, configure multiple entries, one for the 0 offset and one for the 0x81 offset, to account for the different ways in which the two operating systems assign base MAC addresses.

# Feature-Based Licensing

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

**Table 24: License models**

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

## Supported Browsers

Use the following browser versions to access Enterprise Device Manager (EDM):

- Microsoft Edge 97
- Microsoft Internet Explorer 11+
- Mozilla Firefox 96
- Google Chrome 97
- Safari 15.3

➡ **Important**
For optimal performance, use Mozilla Firefox or Google Chrome.

## Feature Differences

Extreme Networks has implemented feature parity between the platforms with a few exceptions. Some features are supported on one platform and not another to maintain compatibility with previous releases. In other cases, the difference reflects the differing roles of the switches in the network..

For information about feature support across all platforms, see *Fabric Engine Feature Support Matrix*.

## Memory Usage

These switches intentionally reboot when memory usage on the switch reaches 95%.
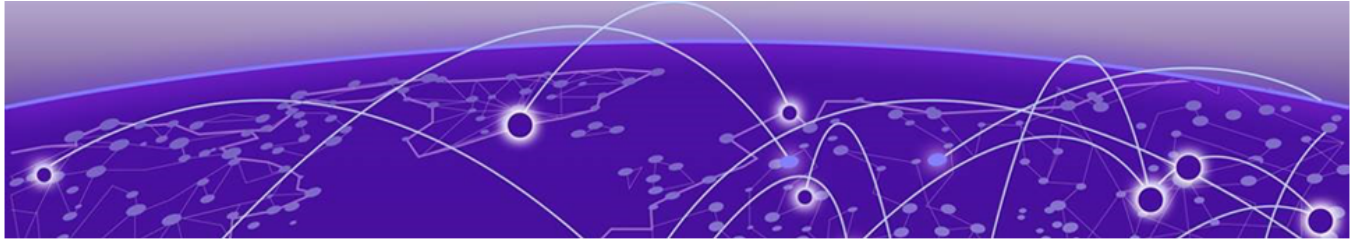
# Known Issues and Restrictions

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

## Known Issues

This section identifies the known issues in this release.

### Known Issues for 8.6

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

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-1280 | The following error message occurs when performing shutdown/no-shutdown commands continuously: `IO1 [05/02/14 06:59:55.178:UTC] 0x0011c525 00000000 GlobalRouter COP-SW ERROR vsp4kTxEnable Error changing TX disable for SFP module: 24, code: -8` | None. When this issue occurs, the port in question can go down, then performs a shutdown/no-shutdown of the port to bring it up and resumes operation. |
| VOSS-1285 | CAKs are not cleared after setting the device to factory-default. | None. Currently this is the default behavior and does not affect functionality of the MACsec feature. |
| VOSS-1288 | Shutting down the T1 link from one end of the link does not shut down the link at the remote end. You could experience traffic loss if the remote side of the link is not shut down. | This issue occurs only when a T1 SFP link from one end is shutdown. Enable a dynamic link layer protocol such as LACP or VLACP on both ends to shut the remote end down too. As an alternative, administratively disable both ends of the T1 SFP link to avoid the impact. |
| VOSS-1289 | On a MACsec-enabled port, you can see delayed packets when the MACsec port is kept running for more than 12 hours. This delayed packet counter can also increment when there is complete reordering of packets so that the application might receive a slow response. But in this second case, it is a marginal increase in the packet count, which occurs due to PN mismatch sometimes only during Key expiry, and does not induce any latency. | None. |
| VOSS-1309 | You cannot use EDM to issue **ping** or **traceroute** commands for IPv6 addresses. | Use CLI to initiate **ping** and **traceroute** commands. |
| VOSS-1310 | You cannot use EDM to issue **ping** or **traceroute** commands for IPv4 addresses. | Use CLI to initiate **ping** and **traceroute** commands. |
| VOSS-1312 | On the VSP 8400 Series 40-gigabit ports, the small metallic fingers that surround the ports are fragile and can bend out of shape during removal and insertion of the transceivers. When the fingers are bent, they prevent the insertion of the QSFP+ transceiver. | Insert the QSFP+ carefully. If the port becomes damaged, it needs to be repaired. |

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

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

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

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

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-2033 | The following error messages appear when you use the **shutdown** and **no shutdown** commands on the MLT interface with ECMP and BGP+ enabled:<br>`CP1 [01/23/16 11:10:16.474:UTC] 0x00108628 00000000 GlobalRouter RCIP6 ERROR rcIpReplaceRouteNotifyIpv6:FAIL ReplaceTunnelRec conn_id 2`<br>`CP1 [12/09/15 12:27:02.203:UTC] 0x00108649 00000000 GlobalRouter RCIP6 ERROR ifyRpcOutDelFibEntry: del FIB of Ipv6Route failed with 0: ipv6addr: 201:6:604:0:0:0:0:0, mask: 96, nh: 0:0:0:0:0:0:0:0 cid 6657 owner BGP`<br>`CP1 [12/09/15 12:20:30.302:UTC] 0x00108649 00000000 GlobalRouter RCIP6 ERROR ifyRpcOutDelFibEntry: del FIB of Ipv6Route failed with 0: ipv6addr: 210:6:782:0:0:0:0:0, mask: 96, nh: fe80:0:0:0:b2ad:aaff:fe55:5088 cid 2361 owner OSPF` | Disable the alternate path. |
| VOSS-2036 | IPsec statistics for the management interface do not increment for inESPFailures or InAHFailures. | None. |
| VOSS-2117 | If you configure static IGMP receivers on an IGMPv3 interface and a dynamic join and leave are received on that device from the same destination VLAN or egress point, the device stops forwarding traffic to the static receiver group after the dynamic leave is processed on the device. The end result is that the IGMP static groups still exist on the device but traffic is not forwarded. | Disable and re-enable IGMP Snooping on the interface. |
| VOSS-2128 | EAP Security and Authentication EDM tabs display additional information with internal values populated, which is not useful for the end user. | There is no functional impact. Ignore the additional information in EDM. Use the CLI command **show eapol port interface** to see port status. |
| VOSS-2207 | You cannot configure an SMTP server hostname that begins with a digit. The system displays the following error:<br>`Error: Invalid IP Address or Hostname for SMTP server` | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-2208 | While performing CFM Layer 2 traceroute between two BEBs via a transit BCB, the transit BCB hop is not seen, if the transit BCB has ISIS adjacencies over FE l3core with both source BEB and destination BEB. | None. |
| VOSS-2253 | Trace level command does not list module IDs when '**?**' is used. | To get the list of all module IDs, type **trace level,** and then press **Enter.** |
| VOSS-2285 | When on BEB, continuously pinging IPv6 neighbor address using CLI command **ping -s**, ping packets do not drop, but instead return no answer messages. | Restart the ping. Avoid intensive CPU processing. |
| VOSS-2333 | Layer 2 ping to Virtual BMAC (VBMAC) fails, if the VBMAC is reachable via Layer 2 core. | None. |
| VOSS-2418 | When you configure and enable the SLA Mon agent, the SLA Mon server is able to discover it but the agent registration on the switch does not occur. | None. |
| VOSS-2422 | When a BGP Neighbor times out, the following error message occurs: `CP1 [03/11/16 13:43:39.084:EST] 0x000b45f2 00000000 GlobalRouter SW ERROR ip_rtdeleteVrf: orec is NULL!` | There is no functional impact. Ignore the error message. |
| VOSS-2859 | You cannot modify the port membership on a protocol-based VLAN using EDM, after it has been created. | Use CLI to provision the port membership on the protocol-based VLAN or delete the protocol-based VLAN, and then re-create it with the correct port member setting. |
| VOSS-3393 | When the SLA Mon agent IP is created on a CLIP interface, the switch provides the CLIP-id as the agent MAC. | There is no functional impact. Use different CLIP IDs to differentiate the SLA Mon agents from the SLA Mon server. |
| VOSS-4255 | If you run IP traceroute from one end host to another end host with a DvR Leaf in between, an intermediate hop will appear as not responding because the Leaf does not have an IP interface to respond. The IP traceroute to the end host will still work. | None. |

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

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

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-8516 | Secure Copy (SCP) cannot use 2048-bit public DSA keys from Windows. | Use 1024/2048-bit RSA keys or 1024-bit DSA keys. |
| VOSS-9516 | When you connect to EDM using HTTPS, you can see multiple `SSL negotiation with client successful` messages during your EDM session. The system displays this message, each time a successful SSL_Handshake occurs between the web browser and the web server. The log file cannot show as many messages as the console and the timing between messages can be different because logging does not occur in real time. | None. |
| VOSS-9589 | Dynamic Nickname Assignment is not supported over Fabric Extend tunnels. | None. |
| VOSS-9621 | On these products, 1G Copper Pluggable auto-negotiation is always enabled after a reboot, despite configuration settings. | If you do not want to use auto-negotiation, disable it after the reboot. |
| VOSS-9921 | Bootup redirection timeout is longer than the UNI port (SMLT) unlock timer. If both vIST nodes boot together in factory default configuration fabric mode or without a nickname, the vIST ports will not enable for up to 4 minutes. During the delay the nickname server is unreachable and vIST is not online. | None. |
| VOSS-10380 | If you enable and configure IPv6 Source Guard and EAPoL on a port, and create and configure a Guest VLAN on the same port without DHCP Snooping and ND-inspection, no error is shown. The port is not added to the Guest VLAN. | None. |
| VOSS-10381 | If you enable and configure IPv6 Source Guard and EAPoL MHSA on a port, and create and configure RAVs for Non-EAP clients on the same port without DHCP Snooping and ND-inspection, no error is shown. The client displays as authenticated into RAV, even when port is not a member of RAV. | None. |
| VOSS-10412 | Removal of the QSFP+ to SFP+ adapter with a 10G pluggable is not detected on the VSP 8404 and VSP 8404C when in non channelized mode. | The QSFP+ to SFP+ adapter and detection works only on ports with channelization enabled. |
| VOSS-10574 | IS-IS sys-name output is not truncated for **show isis spbm nick-name** or **show ip route** commands. If a long character sys-name is in use, the full sys-name display can cause misalignment of the output columns. | None. |

| 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-10891 | DvR leaf vIST: Wrong rarSmltCheckSmltPeerMac MLT warning displays when the peer vIST MAC address is learned from local | None. rarSmltCheckSmltPeerMac MLT warning has no functional impact. You can ignore the error message. |
| VOSS-11895 | In a vIST SMLT environment where streams are both local and remote, if source and receiver port links are removed and reinserted several times, eventually traffic will not be forwarded to local single-homed receivers on one peer if the traffic is ingressing from the vIST peer over the NNI link. If the stream ingresses locally, it is received by the local UNI receivers. | Disable and re-enable Fabric Multicast (`spbm <1-100> multicast enable`) on the source VLAN to be able to delete the streams and come back in properly. |
| VOSS-11943 | This release does not support per-port configuration of Application Telemetry. Because the feature is enabled globally and VSP 7432CQ supports 32 100 Gbps ports, an undesirable condition could be encountered when an exceeded amount of Application Telemetry mirrored packets are sent to the collector. | None. |
| VOSS-12330 | When accessing the on-switch RESTCONF API documentation in a web browser, the page does not render correctly. | Ensure you include the trailing slash (/) in the URL: `http(s)://<ip-address>:8080/apps/restconfdoc/`. For more information, see *Fabric Engine User Guide*. |
| VOSS-12405 | To reach a VM, all front panel traffic must travel through an Insight port, which is a 10 Gbps port. If front panel port traffic is over 10 Gbps, this situation represents an over subscription on the Insight port and some of the packets will be dropped. As a result, Extreme Management Center or ExtremeCloud IQ - Site Engine can lose connectivity to the Analytics engine if Application Telemetry is enabled. | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-13159 | The ixgbevf Ethernet device driver within the TPVM does not correctly handle the interface MTU setting. Specifically, if you configure the interface in SR-IOV mode, packets larger than the MTU size are allowed. | To avoid this problem, configure the desired MTU size on both the relevant front-panel port and Insight port from the NOS CLI. |
| VOSS-13463 | Out port statistics for MLT port interfaces are not accurate. | Use the command `show io nic-counters` to display detailed port stats and error info on XA1400 Series. |
| VOSS-13667 | An intermittent issue in SMLT environments, where ARPs or IPv6 neighbors are resolved with delay can cause a transient traffic loss for the affected IPv6 neighbors. The situation auto-corrects. | None. |
| VOSS-13680 | Interface error statistics display is inaccurate in certain scenarios. | Use the command `show io nic-counters` to display detailed port stats and error info on XA1400 Series. |
| VOSS-13681 | QoS: **show qos cosq-stats cpu-port** command output is not supported. | Use the command `show io cpu-cosq-counters` to display detailed cosq-stats on XA1400 Series. |
| VOSS-13693 | QoS: Traffic can egress out of the queue at a different ratio than the default configuration. After the guaranteed traffic rate is served to all egress port queues, any excess bandwidth is shared equally to all queues instead of distributing on weight assigned to each queue. | None. |
| VOSS-13702 | Do not use the ACE actions of *deny* and *mirror-to-isid* together on VSP 7400 Series. | None. |
| VOSS-13717 VOSS-14393 VOSS-14972 | Link on remote side doesn't go down after admin shut on XA1400 while using 10G DAC or a 4x10 - 40 G breakout DAC. On the XA1400 side link goes down but Link LED shows as up. Both 10G and 4x10G DAC are not fully supported because of this issue | None for DAC and breakout cables. Because of this issue, the following optical transceivers are not supported:<br>• AA1404036-E6<br>• AA1404042-E6<br>• C9799X4-5M |
| VOSS-13794 | You cannot use SFTP to transfer files larger than 2 GB to the switch. | Use SCP. |

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

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

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-15313 | On a VSP 8404C switch using an 8424XT ESM, on a link with MACsec connectivity on both ends, and Auto-Negotiation advertisements set to 10000-full, the link will not come back up if the ESM is hot-swapped or the slot is reset. | To avoid this issue, disable MACsec prior to the hot swap or reset, and then re-enable.<br>If you have experienced the issue, shut either one of the link ports down and bring it back up. |
| VOSS-15391 | An SNMP walk on the **rcIgmpSnoopTraceTable** table will fail with an **OID not increasing** error. CLI and EDM are unaffected by this issue. | None. |
| VOSS-15463 | XA1440 and XA1480 switches can experience intermittent Link Up and Link Down transitions on the 10/100/1000BASE-T Ethernet ports upon booting. | No workaround, but there is no functional impact. |
| VOSS-15541 | You can experience temporary traffic loss when shutting down an LACP SMLT port (and therefore causing the local SMLT to go down), in a network with scaled Multicast traffic over an SPB cloud, while the datapath processes all dpm letter messages during LCAP recovery. This slow LACP recovery situation is only seen with scaled Multicast traffic over an SPB cloud. | Use static MLTs. |
| VOSS-15605 | When you delete the VLAN IDs from the assigned I-SID of two vIST peers, the second VLAN ID deletion triggers log report **0x0013851e** from the first peer, indicating that a Layer 3 MAC address deletion has failed. | No workaround, but there is no functional impact—the MAC address was deleted when the VLAN:ISID association was deleted. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-15812 | Layer 3VSN IPv4 BGP (and static) routes having their next-hops resolved via IS-IS routes could result in traffic loss. | Choose the following workarounds, based on your deployment and needs:<br>• Use static routes to reach the loopbacks used as BGP peers, (static routes having better preference than IS-IS); use static routes with next-hops reachable on the UNI side (L2VSN).<br>• Use OSPF to reach the loopbacks used as BGP peers, but take care to ensure that the OSPF route towards the BGP peer is chosen as the "best route" (as IS-IS has a better preference than OSPF). There are several ways to accomplish this—either don't redistribute that route in IS-IS if it is not needed, or control the redistribution with a route-map, etc.<br>• Have BGP peers reachable directly via a C-VLAN; do not use loopback interfaces as BGP peer addresses.<br>• If none of the above workaround scenarios are suitable for your deployment, do not use internal Border Gateway Protocol (iBGP) peering. |
| VOSS-15878 | VSP 4900 Series and VSP 7400 Series do not boot with just the serial console cable connected and no terminating device, for example, a terminal server, PC, or Mac. | Either attach terminal equipment or disconnect the console cable. |
| VOSS-16221 | Layer 2 ping does not work for packets larger than 1300 on an XA1400 Series. | Use Layer 2 ping with packets smaller than 1300 bytes. |
| VOSS-16365 | Running the command **show pluggable-optical-module detail** on an XA1400 Series device is highly CPU intensive to read and reply with the EEPROM details. Due to a delay in ethtool response, a watchdog miss event can occur and the event is recorded in the /intflash/wd_stats/1/wd_stats.ssio.1.log file. This scenario occurs more often if 10Gb SFP+ optics with DDM capability are installed. | None. The high CPU usage and response delay for this command is expected and cannot be resolved. No console log is generated. When the scenario occurs, the Watchdog outage is approximately 5 seconds. |
| VOSS-16436 | Using the console connection on an XA1400 Series device while running a show command with large data output can result in drops of processing control packets. | Use Telnet or SSH connectivity instead of console connection. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-16951 | On a VSP4900-48P, VSP4900-24S and VSP 7400 Series devices, if you run the **show boot config sio** CLI command before you have configured the baud rate, the output of the command is empty. | Configure the baud rate before you run the **show boot config sio** command. The only supported baud rate for these devices is 115200. |
| VOSS-16971 | On VSP4900-24S, VSP4900-24XE, 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-17002 | For ingress packets that are larger than the system MTU size on XA1400 Series ports 1/1 through 1/4, error counters do not increment in the **show interfaces gigabitethernet error** CLI command. | Use the **show io nic-counters** CLI command to verify if the tx_error counters are getting incremented. If they are getting incremented, the packets are getting dropped at egress. If they are not getting incremented, the packets are getting forwarded. |
| VOSS-17429 | For XA1400 Series devices connected to an FE tunnel over IPsec in a dual NAT scenario, if the IPsec responder is rebooted continuously multiple times, the tunnel cannot come back up. | Manually disable and then re-enable IPsec under the Initiator's ISIS logical interface. |
| VOSS-17478 | On 1 G-capable VSP 4900 Series devices, the platform MACsec statistics cannot match the port Interface statistics after Key expiry. | No Workaround. This is a Statistics data issue where the expired SA Packets Counts are removed and not accounted. There is no packet loss, and no errors. |
| VOSS-17523 | If an FE tunnel goes down between two connected XA1400 Series devices, an MTU Warning console message is logged if a ping request is issued while the tunnel is down. | You can safely ignore this warning message. |
| VOSS-17567 | Do not use the inter-vrf /32 static routes defined with a next-hop IP address that resides in a different destination next-hop-vrf context. | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-18023 | The management port on the 5520 switch does not support Auto-MDIX (the automatic detection of transmit and received twisted pairs).<br>As a best practice, enable the default auto-negotiation setting on the management port.<br>Because the management port does not support Auto-MDIX, when auto-negotiation is disabled, a crossover cable might be necessary to have the port link up and pass traffic.<br><br>**Note:** If the peer device supports Auto-MDIX, then either a straight through or crossover will work. The issue occurs only if both ends of the connection do not support Auto-MDIX. | None. |
| VOSS-18238 | When a management VLAN with DHCP is used to reach a RADIUS server, and the RADIUS server cannot be reached, the system waits for 15 minutes before attempting to reach the RADIUS server again. This is true even if the RADIUS server becomes reachable before the 15 minutes have elapsed. | None. |
| VOSS-18278 | On the 5520 switch, when you make any change relating to port speed, the port statistics are cleared. This is applies to all front panel fiber and copper ports as well as VIM ports.<br>The following are examples of changes relating to port speed:<br>• Changing the auto-negotiation configuration settings on a copper port<br>• Different negotiated speed on a copper port<br>• Changing out an optical device for one having a different speed, for example changing from 1 Gb to 10 Gb | None. |
| VOSS-18360 | This is an intermittent issue on the VSP 7400 Series with no impact to functionality, ISIS is disabled while the **show fulltech** command is running on a telnet session. Due to this the fulltech command will not find the expected I-SID value, as it is removed by the **no isis** command. | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-18477 | On the VSP 4900 Series, an intermittent traffic loss over the FE tunnels, in SMLT contexts, occurs for a few seconds, when you read ports to the SMLT trunk. | None. |
| VOSS-19212 | After upgrading a VSP 7432CQ switch to VOSS 8.2.5 and rebooting, the presence of a faulty power supply unit will cause the system to terminate. A message in the debug log will report that the software could not read the contents of the power supply's EEPROM (*carbonatelib_ps_read_eeprom* operation). | Replace the power supply unit in the switch. |
| VOSS-19253 | On 5520 switches, authentication is not allowed for requesters that use the switch's MAC address as destination rather than using the 802.1x reserved MAC address. | None. |
| VOSS-19260 | Port mirroring does not work on port 1/s1 of VSP 7400-48Y if the connection type is OVS/SR-IOV. | Use a connection type of VT-d for port 1/s1. |
| VOSS-19364 | If you use a Windows client to create a direct SSH session with the Fabric IPsec Gateway virtual machine (VM) and make configuration changes, the VM does not display the running configuration after the SSH session drops. This issue does not affect existing configuration or traffic; only new configuration or show commands are affected. | Use the **reboot** command to reboot the VM. The VM loads the configuration as normal but services are temporarily impacted during the reboot. |
| VOSS-19827 | LLDP IPv6 neighbors do not display in EDM. LLDP IPv6 is only supported in CLI. | To display LLDP IPv6 neighbors, use the **show lldp neighbor summary** command. |
| VOSS-20115 | You cannot change the management VLAN interface discovered on XA1400 Series in Extreme Management Center or ExtremeCloud IQ - Site Engine as part of Zero Touch Provisioning Plus (ZTP+). XA1400 Series does not support the OOB interface. You can only use the discovered interface and change other configuration values. | On XA1400 Series, use the discovered interface within Extreme Management Center or ExtremeCloud IQ - Site Engine for basic onboarding. Use either Extreme Management Center or ExtremeCloud IQ - Site Engine or CLI to complete the remaining configuration. |
| VOSS-20200 | For VSP 8404C, if you remove and insert an Ethernet Switch Module (ESM), which has NNI ports that are members in an LACP-dynamic MLT, some ports are intermittently missing in the dynamic MLT after the ESM insertion. Traffic is affected for streams that need to exit the NNI links over the dynamic MLT for the missing ports. Rebooting the switch returns the ports to the dynamic MLT. | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-20227 | On XA1400 Series, the VOSS OS time does not synchronize to the real time clock (RTC) after system reboot. After the switch completely boots, NTP synchronization occurs and the VOSS OS has the correct time. The OS time can be incorrect for up to two minutes after system reboot. | None. |
| 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:<br>• `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`<br>• `1 2021-02-17T23:32:16.811+01 :00 DIST-H9-E3.1-01 CP1 - 0x000a45ac - 00000000 GlobalRouter RADIUS ERROR rad_processPendingRequest: unable to send request` | None. This issue has no functional impact. |
| VOSS-20456 | Although the Management Router is not supported in the NOS, you can add a static route for VRF 512 using EDM. The route does not become active even if the next-hop address is reachable from the OOB management interface. | None. This issue has no functional impact. |
| VOSS-21097 | In Multi-Area where vIST peers are boundary nodes, vIST can briefly flap during connection formation when IS-IS is disabled and then reenabled on both vIST peers. | None. |
| VOSS-21123 | Brouters on UNIs of VSP 7400 vIST peers cannot ping each other. | Add a static ARP for the Brouter of the VIST peer. |
| VOSS-21233 | Clearing DVR host entries in a highly scaled Multi-Area DVR environment may trigger DBSYNC WARNING messages (`0x00390606 - 00000000 GlobalRouter DBSYNC WARNING Message queue length from DB Sync to tMain reached warning threshold`) but these can be expected in a scaled environment and are not a malfunction. | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-21432 | When the **show trace file** command output is saved to a file, non-ASCII characters display at the end of the string. | None. |
| VOSS-21607 | In scaled Multi-area SPB networks, Layer 3 VSN traffic loss can occur on VSP 8600 Series vIST peers from the Remote Area connected to two VSP 7400 Series switches in the Home Area after the remote area virtual node System ID is changed on VSP 7400 Series Boundary Nodes manually (for instance, in this order: IS-IS remote instance is disabled, virtual node system ID is changed to a new value, and IS-IS remote instance is re-enabled). | This issue is not seen if the Virtual node system ID is automatically derived. |
| VOSS-21663 | EAP LLDP authentication (RADIUS bypass) does not work with Cisco IP Phone 7821. With Auto-sense, the authentication of the phone seems to complete successfully and is reachable, but after a while when it is de-authenticated, the phone is also not reachable. This is because Cisco is not encoding its MAC in LLDP packets, but VOSS expects the MAC address in the LLDP packets. | None. |
| VOSS-21842 | Under some circumstances after reboot of a 5520 Series switch with a specific configuration, the SPB nickname fails to be assigned. The specific configuration includes Fabric Attach using a Universal Ethernet port configured as a channelized MLT. | Disable and re-enable IS-IS on the MLT interface. |
| VOSS-21964 | When using Windows SCP application on a switch to transfer a file, an error message displays even if a file transfers successfully. | |
| VOSS-22255 | Ping, which originates from a local CP, fails for ICMP packets bigger than 1500 sent from Layer 3 VSN interface. | Initiate ping with packets size smaller than 1500. |
| VOSS-22264 | Sending traffic from duplicate DvR hosts over a long period of time can cause errors. | Remove duplicate DvR hosts from the network. |
| VOSS-22281 | An ICMP echo packet reply is not sent for jumbo ICMP echo requests between 9568 and 9600 bytes over a MACsec connection to 5420 Series. | Configure MTU to less than 9568 bytes on neighboring devices. |
| VOSS-22522 | RESTCONF is delayed in a scaled setup with 2,000 VLANs. | None. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-22824 | When an IPv6 multicast entry is learned in a PIM over IPv6 environment on the 5420 Series, an error message displays. | None. |
| VOSS-22848 | Host entry displays at the original location when you migrate a DvR host to another domain. | Clear DvR host entries.<br>This also recovers when the affected controller receives another ARP request or Layer 2 packet originated by the host. |
| 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-22917 | You cannot add non-private VLAN ports to a platform VLAN that already has a private-VLAN member. | Add non-private VLAN ports to a platform VLAN before you create a Switched UNI for the isolated or promiscuous port in the platform VLAN. |
| VOSS-23018 | DEI marked packets drop when traffic egresses from an untagged port. The DEI bit is configured when traffic is marked yellow by a policer. The drop happens on the switch that is downstream from the switch with the policer. | Choose one of the following:<br>• Disable ingress policer and port rate limiter functionality.<br>• Do not configure packets as yellow traffic. |
| 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-23205 | If you repeatedly enable and disable Auto-Negotiation, ports can remain in the operationally down state. | 1. Disable and enable the port.<br>2. If the port remains down, toggle the Auto-Negotiation status on the port again.<br>3. If the port remains down, restart the switch. |
| VOSS-23216 | If you do not enable the DvR interface when you configure a dvr-one-ip interface, the dvr-one-ip interface does not display when you issue the **show dvr interfaces** command. | Enable the DvR interface. |
| VOSS-23229 | In an E-Tree scenario, IPv6 packets are forwarded between isolated ports on 5520 Series, 5420 Series, and VSP 7400 Series. | None. |
| VOSS-24742 | A port remains down if you do not synchronize the system clock, or if it fails to synchronize, with the NTP current time. | Ensure that the system clock synchronizes to the current time and toggle the port link state. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-24771 | When you configure the `macsec connectivity-association name` to the maximum of 16 characters using CLI, the connectivity association name attached to the port or interface does not display in EDM. | Configure the `macsec connectivity-association name` to a maximum of 15 characters in CLI. |
| VOSS-24777 | In the following port configurations on 5520 Series, 5420 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, and VSP 8400 Series inVSN ACL entries match ingressing packets that have the same VID as the VLAN associated with the ACL I-SID even if the ACL inVSN I-SID is different:<br>• on an S-UNI port without a platform VLAN<br>• on a T-UNI port VLAN | None. |
| VOSS-24786 | The 24-port models of 5320 Series do not always reach the advertised MACsec bandwidth. | None. |
| VOSS-24872 | If the collector reachability path changes for Application Telemetry, it is not reflected properly in CLI. Packets remain mirrored towards the correct path but CLI does not reflect the next hop. | None. There is no functional impact. |
| VOSS-25001 | When you configure IPv6 on the same DvR-enabled VLAN, IPv6 connectivity might be lost between vIST peers when you add an MLT in that VLAN. | None. |
| VOSS-25063 | In EDM (**Configuration** > **Edit** > **Chassis**) the Type for Fabric Engine platforms contains VOSS. The same issue is seen if you perform an SNMP Get on the rcChasType MIB object. | None. |
| VOSS-25078 | MAC addresses learned on a Switched UNI (S-UNI) port cannot be flushed. | None. |
| VOSS-25023 | 5520 Series, 5420 Series, and 5320 Series platforms can reach 100% CPU utilization during inband transfer (FTP, SFTP, and SCP). | None. |
| VOSS-25158 | If you disable and re-enable the operational state of an Auto-sense enabled port, STP is enabled when the interface is in the Auto-sense WAIT state; STP should be disabled in all Auto-sense port states. | Disable and re-enable Auto-sense on the interface. |

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-25162 | RESTCONF ARP and MAC data: on 5x20 switches with 5K ARP entries and 5K MAC entries, it takes approximately 1 minute to retrieve data. The time increases based on the number of entries.<br>The same occurs on VSP 7400 Series with over 15K entries. | None. |
| VOSS-25225 | On 5320 Series, the four highest SFP+ ports are available at 10 Gbps with Trial Licenses. After license expiration, the port speeds drop to 1 Gbps. | Use the **extend-time-period** command prior to the expiration of the Trial License. |
| VOSS-25229 | The SNMP server community credentials need to be recreated after a power cycle. | None. To avoid this situation, use the following steps to properly power down and restart:<br>1. Use the **sys shutdown** command to power down the switch.<br>2. Physically disconnect the power cord and wait 3 to 5 minutes before restoring power to the chassis. |

# Restrictions and Expected Behaviors

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

For Port Mirroring considerations and restrictions, see *Fabric Engine User Guide*.

## General Restrictions and Expected Behaviors

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

**Table 25: 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 25: General restrictions (continued)**

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-687 | EDM and CLI show different local preference values for a BGP IPv6 route.<br>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.<br>CLI displays path attributes associated with the route entry, which can be modified by a policy. If a route policy is not configured, the local preference shows the default value of 100. | None. |
| VOSS-1954 | After you log in to EDM, if you try to refresh the page by clicking on the refresh button in the browser toolbar, it will redirect to a blank page. This issue happens only for the very first attempt and only in Firefox. | To refresh the page and avoid this issue, use the EDM refresh button instead of the browser refresh button. If you do encounter this issue, place your cursor in the address bar of the browser, and press **Enter**. This will return you to the EDM home page. |
| VOSS-2166 | The IPsec security association (SA) configuration has a NULL Encryption option under the **Encrpt-algo** parameter. Currently, you must fill the **encrptKey** and **keyLength** sub-parameters to set this option; however, these values are not used for actual IPsec processing as it is a NULL encryption option. The NULL option is required to interoperate with other vendors whose IPsec solution only supports that mode for encryption. | There is no functional impact due to this configuration and it only leads to an unnecessary configuration step. No workaround required. |
| VOSS-21946 | When you create a vrf using the POSTMAN API platform, special characters, such as \\\\ and ### included in the URL are ignored. | None. |

**Table 25: General restrictions (continued)**

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

**Table 25: General restrictions (continued)**

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-10168 | The system CLI does not prevent you from using the same IP address for the VXLAN Gateway hardware VTEP replication remote peer IP and OOB Management IP. | Manually check the IP configured as the OOB Management IP. Do not use the OOB Management IP address as the replication remote peer IP address. |
| VOSS-11817 | The OVS connect-type for virtual service Vports is designed in such a way that it connects to any generic virtual machine (VM) guest OS version using readily available Ethernet device drivers. This design approach provides initial connectivity to the VM in a consistent manner. A consequence of this approach is that Vports created with connect-type OVS will show up as 1 Gbps interfaces in the VM even though the underlying Ethernet connection supports 10 Gbps . | If additional performance is desired, upgrade the VM guest OS with an Ethernet device driver that supports 10 Gbps interfaces. |
| VOSS-12151 | If logical switch has only hardware ports binding, and not VM behind software VTEP, Broadcast, Unknown Unicast, and Multicast (BUM) traffic does not flow between host behind two hardware VTEP. The NSX replicator node handles the BUM traffic. NSX does not create the replicator node unless a VM is present. In an OVSDB topology, it is expected that at least one VM connects to the software VTEP. This issue is an NSX-imposed limitation. | After you connect the VM to the software VTEP, the issue is not seen. |
| VOSS-12395 | You cannot use the following cables on 10 Gb fiber interfaces, or 40 Gb channelized interfaces, with the QSA28 adapter:<br>• 1, 3, and 5 meter QSFP28 25 Gb DAC<br>• 20 meter QSFP28 25 Gb AOC | n/a |
| VOSS-17871 | Starting with VOSS 8.1.5, internal system updates have resulted in a more accurate accounting of memory utilization. This can result in a higher baseline memory utilization reported although actual memory usage is not impacted. | Update any network management alarms that are triggered by value with the new baseline. |

**Table 25: General restrictions (continued)**

| Issue number | Description | Workaround |
|---|---|---|
| VOSS-18523 | When you configure a port using Zero Touch Provisioning Plus (ZTP+) with Extreme Management Center or ExtremeCloud IQ - Site Engine, the port cannot be part of both a tagged VLAN and an untagged VLAN. | n/a |
| VOSS-18409 | On the XA1400 Series switches, only one Central Processing Unit (CPU) core is assigned for control plane protocol processing. In a highly scaled scenario, a port toggling or negative scenario keeps the CPU core busy in updating the software datapath entries. Similarly, some show CLI commands that require a lot of data gathering keep the CPU core busy. In such a scenario, the main task which is responsible for handling protocol packets like Bidirectional Forwarding Detection, Intermediate-System-to-Intermediate-System, Virtual Link Aggregation Control Protocol, and so on is busy. | For scaled scenarios on XA1400 Series switches, the CLI commands that have large sections of output, for example, show fulltech, show io spb tables, and show tech, the output must be redirected into a file. |
| VOSS-18774 | SSL negotiation fails when using OpenSSL client version 1.1.1. With OpenSSL 1.1.1, the server-name extension is used. This extension needs to equal the domain name in the server certificate, otherwise the certificate lookup on the server fails because the FIPS 140-2 certified cryptographic module processes the server-name extension. | Can connect using: `bash# openssl s_client -connect <domain-name>:443` |
| VOSS-18851 | Do not define a static route in which the NextHop definition uses an Inter-VRF redistributed route. Such a definition would require the system to perform a double lookup. When you attempt to define a static route in this way, an error message is generated. | Define the static route in such a way that it does not require Inter-VRF redistributed routing. |

Table 25: 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 |
| 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 |

**Table 25: General restrictions (continued)**

| Issue number | Description | Workaround |
|---|---|---|
| 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:<br>• On an IGMP snoop-enabled interface, you can flush IGMP sender records.<br>**Caution:**<br>Flushing sender records can cause a transient traffic loss.<br>• On an IGMP-enabled Layer 3 interface, you can toggle the IGMP state.<br>**Caution:**<br>Expect traffic loss until IGMP records are built after toggling the IGMP state. |
| wi01145099 | IP multicast packets with a time-to-live (TTL) equal to 1 are not switched across the SPB cloud over a Layer 2 VSN. They are dropped by the ingress BEB. | To prevent IP multicast packets from being dropped, configure multicast senders to send traffic with TTL greater than 1. |
| wi01159075 | VSP 4450GTX-HT-PWR+: Mirroring functionality is not working for RSTP BPDUs. | None. |
| wi01171670 | Telnet packets get encrypted on MACsec-enabled ports. | None. |
| wi01198872 | On VSP 4450 Series, a loss of learned MAC addresses occurs in a vIST setup beyond 10k addresses.<br>In a SPB setup the MAC learning is limited to 13k MAC addresses, due to the limitation of the internal architecture when using SPB. Moreover, as vIST uses SPB and due to the way vIST synchronizes MAC addresses with a vIST pair, the MAC learning in a vIST setup is limited to 10K Mac addresses. | None. |
| wi01210217 | The command **show eapol auth-stats** displays LAST-SRC-MAC for NEAP sessions incorrectly. | n/a |

**Table 25: General restrictions (continued)**

| Issue number | Description | Workaround |
|---|---|---|
| wi01211415 | In addition to the fan modules, each power supply also has a fan. The power supply stops working if a power supply fan fails, but there is no LED or software warning that indicates this failure. | Try to recover the power supply fan by resetting the switch. If the fan does not recover, then replace the faulty power supply. |
| wi01212034 | When you disable EAPoL globally:<br>• Traffic is allowed for static MAC configured on EAPoL enabled port without authentication.<br>• Static MAC config added for authenticated NEAP client is lost. | n/a |
| wi01212247 | BGP tends to have many routes. Frequent additions or deletions impact network connectivity. To prevent frequent additions or deletions, reflected routes are not withdrawn from client 2 even though they are withdrawn from client 1. Disabling route-reflection can create a black hole in the network. | Bounce the BGP protocol globally. |
| wi01212585 | LED blinking in EDM is representative of, but not identical to, the actual LED blinking rates on the switch. | n/a |
| wi01213040 | When you disable auto-negotiation on both sides, the 10 Gbps copper link does not come up. | n/a |
| wi01213066 wi01213374 | EAP and NEAP are not supported on brouter ports. | n/a |
| wi01213336 | When you configure tx mode port mirroring on T-UNI and SPBM NNI ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports appear on the mirror destination port, although they do not egress the mirror source port. This is because tx mode port mirroring happens on the mirror source port before the source port squelching logic drops the packets at the egress port. | n/a |
| wi01219658 | The command **show khi port-statistics** does not display the count for NNI ingress control packets going to the CP. | n/a |
| wi01219295 | SPBM QOS: Egress UNI port does not follow port QOS with ingress NNI port and Mac-in-Mac incoming packets. | n/a |

**Table 25: General restrictions (continued)**

| Issue number | Description | Workaround |
|---|---|---|
| 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:<br>• Enable PIM on the edge.<br>• Ensure that IST peers are either RP or DR but not both. |
| wi01224683<br>wi01224689 | Additional link bounce can occur on 10 Gbps ports when toggling links or during cable re-insertion.<br>Additional link bounce can occur with 40 Gbps optical cables and 40 Gbps break-out cables, when toggling links or during cable re-insertion. | n/a |
| wi01229417 | Origination and termination of IPv6 6-in-4 tunnel is not supported on a node with vIST enabled. | None. |
| wi01232578 | When SSH keyboard-interactive-auth mode is enabled, the server generates the password prompt to be displayed and sends it to the SSH client. The server always sends an expanded format of the IPv6 address. When SSH keyboard-interactive-auth mode is disabled and password-auth is enabled, the client itself generates the password prompt, and it displays the IPv6 address format used in the `ssh` command. | None. |
| wi01234289 | HTTP management of the ONA is not supported when it is deployed with a VSP 4450 Series device. | None. |

## Filter Restrictions

The following table identifies known restrictions.

**Table 26: 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 egress QoS ACL/filters are not supported. |
| All platforms | Control packet action is not supported on InVSN Filter or IPv6 filters generally. |
| All platforms | IPv4/IPv6 VLAN based ACL filters will be applied on traffic received on all the ports if it matches VLAN ID associated with the ACL. |

Table 26: ACL restrictions (continued)

| Applies To | Restriction |
|---|---|
| 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 27: ACE restrictions

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

# Resolved Issues this Release

This release incorporates all fixes from prior releases, up to and including VOSS 8.4.3.1 and VOSS 8.5.0.0.

| Issue number | Description |
| --- | --- |
| VOSS-22325 | Around 1/3rd of I-SID RSPAN mirrored traffic not reaching monitoring server. |
| VOSS-22413 | XA1480 Aggregation box rebooted for unknown reason. |
| VOSS-22859 | XA1400 Series - 8.3.1.0 - slow memory depletion of physical memory and increased use of RSS memory in SSIO since IPSec was enabled. |
| VOSS-22908 | 60 percent of fbufs are in use causing IS-IS adj to drop. |
| VOSS-23043 | 5520 Series does not filter traffic sent with CRC errors. |
| VOSS-23241 | Configured static IP routes do not clear after you configure a switch as a DvR Leaf node. |
| VOSS-24745 | 5520 8.2.8; Connectivity loss in user VLAN when user VLAN co-exist with mgmt VLAN with no IP assigned to mgmt VLAN manually or via mgmt dhcp-client. |
| VOSS-24839 | 8.4.2.1 – Printer cannot access network when directly connected to auto-sense onboarding VSP setup port. |
| VOSS-24841 | VOSS doesn't ignore leading blanks in answers to "y/n" questions. |
| VOSS-24916 | 5520 Series: Inconsistent Auto-Negotiation values on 10G SFP+ connected VIM Ports. |
| VOSS-24925 | Port goes to TRUSTED by using the UNTRUSTED command. |
| VOSS-25069 | Multicast over Fabric Extend over multi-area not working consistently. |

# Related Information

## MIB Changes

### Deprecated MIBs

**Table 28: Common**

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

### Modified MIBs

**Table 29: Common**

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| rc2kPowerConsumptionInfoCardDescription | 1.3.6.1.4.1.2272.1.100.17.1.7 | 8.4 | CHANGE_RANGE: Changed the range from 0..20 to 0..24 |
| rcDvrGlobalRole | 1.3.6.1.4.1.2272.1.219.1.2 | 8.4 | OTHER: parameter controller is not supported on 5420 |
| rcPortEntry | 1.3.6.1.4.1.2272.1.4.10.1.1 | 8.6 | Updated description with difference between GRT and VRF query. |

**Table 29: Common (continued)**

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| rc2kPowerConsumptionInfoCardDescription | 1.3.6.1.4.1.2272.1.100.17.1.7 | 8.6 | CHANGE_RANGE: Changed the range from 0..24 to 0..32 |
| rcAutoSenseIsisHelloAuthKeyId | 1.3.6.1.4.1.2272.1.231.1.1.1.12 | 8.6 | CHANGE_RANGE: Changed the range from 1..255 to 0..255 |

**Table 30: 5420 Series**

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| rcVossSystemCardLedStatus | 1.3.6.1.4.1.2272.1.101.1.1.5.1.4 | 8.4 | ADD_NEW_VALUE: blueBlinking(8) |
| rc2kChassisPortLed3Status | 1.3.6.1.4.1.2272.1.100.1.14 | 8.4 | OTHER: Add 5420 in description, it behaves as 5520. |
| rcVlanNlbMode | 1.3.6.1.4.1.2272.1.3.2.1.46 | 8.4 | NOT_SUPPORT: nlb feature unsupported on platform |
| rcVlanPimGatewayEnable | 1.3.6.1.4.1.2272.1.3.2.1.73 | 8.4 | NOT_SUPPORT: pim gateway feature unsupported on platform |
| rcIpConfPimGatewayEnable | 1.3.6.1.4.1.2272.1.8.1.1.1.30 | 8.4 | NOT_SUPPORT: pim gateway feature unsupported on platform |
| rcIpSpbPimGw | 1.3.6.1.4.1.2272.1.8.114 | 8.4 | NOT_SUPPORT: spb pim gateway feature unsupported on this platform |
| rcIsisPlsbMcastSpbPimGwControllerEnable | 1.3.6.1.4.1.2272.1.63.4.1.15 | 8.4 | NOT_SUPPORT: pim gateway feature unsupported on platform |
| rcIsisPlsbMcastSpbPimGwGatewayEnable | 1.3.6.1.4.1.2272.1.63.4.1.16 | 8.4 | NOT_SUPPORT: pim gateway feature unsupported on platform |
| rcDvrBackboneEntriesAdvControllerName | 1.3.6.1.4.1.2272.1.219.8.1.8 | 8.4 | NOT_SUPPORT: dvr controller feature unsupported on platform |
| rcDvrBackboneEntriesAdvController | 1.3.6.1.4.1.2272.1.219.8.1.9 | 8.4 | NOT_SUPPORT: dvr controller feature unsupported on platform |
| rcPortAutoNegAd | 1.3.6.1.4.1.2272.1.4.10.1.1.62 | 8.5 | ADD_NEW_VALUE: advertise25000Full(13) |

**Table 30: 5420 Series (continued)**

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| SnpxChassisType | | 8.6 | ADD ENUM: m532048T8XE, m532048P8XE, m532024T8XE, m532024P8XE, m532016P4XE, m532016P4XEDC<br>OTHER: Replace "Virtual Services" with "Extreme Networks Fabric Engine" and "VOSS" with "FabricEngine" in comments only for Universal Hardware<br>OTHER: Rebranding for Universal Hardware: Change enum values from m552048TVOSS, m552048WVOSS, m552012MW36WVOSS, m552024TVOSS, m552024WVOSS, m552024XVOSS, m552048SEVOSS to m552048T, m552048W, m552012MW36W, m552024T, m552024W, m552024X, m552048SE |
| rcSysLocatorLED | 1.3.6.1.4.1.2272.1.1.125 | 8.6 | OTHER: Add 5520, 5420 and 5320 in description |
| rcChasType | 1.3.6.1.4.1.2272.1.4.1 | 8.6 | ADD ENUM: a532048T8XEVOSS, a532048P8XEVOSS, a532024T8XEVOSS, a532024P8XEVOSS, a532016P4XEVOSS, a532016P4XEDCVOSS |
| rcIpConfGlobalTcpAdjustMssEnable | 1.3.6.1.4.1.2272.1.8.1.6.29 | 8.6 | OTHER: Add 5320 in description |
| rcIpConfGlobalTcpAdjustMssStatus | 1.3.6.1.4.1.2272.1.8.1.6.30 | 8.6 | OTHER: Add 5320 in description |
| rcIpConfGlobalTcpAdjustMssType | 1.3.6.1.4.1.2272.1.8.1.6.31 | 8.6 | OTHER: Add 5320 in description |
| rcIpConfGlobalTcpAdjustMssValue | 1.3.6.1.4.1.2272.1.8.1.6.32 | 8.6 | OTHER: Add 5320 in description |
| rcIpfixAgingIntervalV2 | 1.3.6.1.4.1.2272.1.66.1.1.5 | 8.6 | OTHER: Add 5320 in description |
| rc2kChassisPortLed3Status | 1.3.6.1.4.1.2272.1.100.1.14 | 8.6 | OTHER: Add 5320 in description |

Table 30: 5420 Series (continued)

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| rc2kBootConfigEnableMacsec | 1.3.6.1.4.1.2272.1.100.5.1.62 | 8.6 | OTHER: Add 5320 in description |
| rc2kCardFrontType | 1.3.6.1.4.1.2272.1.100.6.1.2 | 8.6 | ADD ENUM: voss532048T8XE, voss532048P8XE, voss532024T8XE, voss532024P8XE, voss532016P4XE, voss532016P4XEDC |

Table 31: 5520 Series

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| rcPortAutoNegAd | 1.3.6.1.4.1.2272.1.4.10.1.1.62 | 8.5 | ADD_NEW_VALUE: advertise25000Full(13) |
| SnpxChassisType | | 8.6 | ADD ENUM: m532048T8XE, m532048P8XE, m532024T8XE, m532024P8XE, m532016P4XE, m532016P4XEDC OTHER: Replace "Virtual Services" with "Extreme Networks Fabric Engine" and "VOSS" with "FabricEngine" in comments only for Universal Hardware |
| rcSysLocatorLED | 1.3.6.1.4.1.2272.1.1.125 | 8.6 | OTHER: Add 5520, 5420 and 5320 in description |
| rcChasType | 1.3.6.1.4.1.2272.1.4.1 | 8.6 | ADD ENUM: a532048T8XEVOSS, a532048P8XEVOSS, a532024T8XEVOSS, a532024P8XEVOSS, a532016P4XEVOSS, a532016P4XEDCVOSS |
| rcIpConfGlobalTcpAdjustMssEnable | 1.3.6.1.4.1.2272.1.8.1.6.29 | 8.6 | OTHER: Add 5320 in description |
| rcIpConfGlobalTcpAdjustMssStatus | 1.3.6.1.4.1.2272.1.8.1.6.30 | 8.6 | OTHER: Add 5320 in description |
| rcIpConfGlobalTcpAdjustMssType | 1.3.6.1.4.1.2272.1.8.1.6.31 | 8.6 | OTHER: Add 5320 in description |
| rcIpConfGlobalTcpAdjustMssValue | 1.3.6.1.4.1.2272.1.8.1.6.32 | 8.6 | OTHER: Add 5320 in description |
| rcIpfixAgingIntervalV2 | 1.3.6.1.4.1.2272.1.66.1.1.5 | 8.6 | OTHER: Add 5320 in description |
| rc2kChassisPortLed3Status | 1.3.6.1.4.1.2272.1.100.1.14 | 8.6 | OTHER: Add 5320 in description |

Table 31: 5520 Series (continued)

| Object Name | Object OID | Modified in Release | Modification |
|---|---|---|---|
| rc2kBootConfigEnableMacsec | 1.3.6.1.4.1.2272.1.100.5.1.62 | 8.6 | OTHER: Add 5320 in description |
| rc2kCardFrontType | 1.3.6.1.4.1.2272.1.100.6.1.2 | 8.6 | ADD ENUM: voss532048T8XE, voss532048P8XE, voss532024T8XE, voss532024P8XE, voss532016P4XE, voss532016P4XEDC |

## New MIBs

Table 32: Common

| Object Name | Object OID | New in VOSS Release |
|---|---|---|
| entLogicalTable | 1.3.6.1.2.1.47.1.2.1 | 8.4 |
| entLogicalEntry | 1.3.6.1.2.1.47.1.2.1.1 | 8.4 |
| entLogicalIndex | 1.3.6.1.2.1.47.1.2.1.1.1 | 8.4 |
| entLogicalDescr | 1.3.6.1.2.1.47.1.2.1.1.2 | 8.4 |
| entLogicalType | 1.3.6.1.2.1.47.1.2.1.1.3 | 8.4 |
| entLogicalCommunity | 1.3.6.1.2.1.47.1.2.1.1.4 | 8.4 |
| entLogicalTAddress | 1.3.6.1.2.1.47.1.2.1.1.5 | 8.4 |
| entLogicalTDomain | 1.3.6.1.2.1.47.1.2.1.1.6 | 8.4 |
| entLogicalContextEngineID | 1.3.6.1.2.1.47.1.2.1.1.7 | 8.4 |
| entLogicalContextName | 1.3.6.1.2.1.47.1.2.1.1.8 | 8.4 |
| rcStatOspfIfClearStats | 1.3.6.1.4.1.2272.1.12.5.1.20 | 8.4 |
| ospfRFC1583Compatibility | 1.3.6.1.2.1.14.1.15 | 8.5 |
| ospfOpaqueLsaSupport | 1.3.6.1.2.1.14.1.16 | 8.5 |
| ospfReferenceBandwidth | 1.3.6.1.2.1.14.1.17 | 8.5 |
| ospfRestartSupport | 1.3.6.1.2.1.14.1.18 | 8.5 |
| ospfRestartInterval | 1.3.6.1.2.1.14.1.19 | 8.5 |
| ospfRestartStrictLsaChecking | 1.3.6.1.2.1.14.1.20 | 8.5 |
| ospfRestartStatus | 1.3.6.1.2.1.14.1.21 | 8.5 |
| ospfRestartAge | 1.3.6.1.2.1.14.1.22 | 8.5 |
| ospfRestartExitReason | 1.3.6.1.2.1.14.1.23 | 8.5 |
| ospfAsLsaCount | 1.3.6.1.2.1.14.1.24 | 8.5 |
| ospfAsLsaCksumSum | 1.3.6.1.2.1.14.1.25 | 8.5 |
| ospfStubRouterSupport | 1.3.6.1.2.1.14.1.26 | 8.5 |

Table 32: Common (continued)

| Object Name | Object OID | New in VOSS Release |
|---|---|---|
| ospfStubRouterAdvertisement | 1.3.6.1.2.1.14.1.27 | 8.5 |
| ospfDiscontinuityTime | 1.3.6.1.2.1.14.1.28 | 8.5 |
| spfAreaNssaTranslatorRole | 1.3.6.1.2.1.14.2.1.11 | 8.5 |
| ospfAreaNssaTranslatorState | 1.3.6.1.2.1.14.2.1.12 | 8.5 |
| ospfAreaNssaTranslatorStabilityInterval | 1.3.6.1.2.1.14.2.1.13 | 8.5 |
| ospfAreaNssaTranslatorEvents | 1.3.6.1.2.1.14.2.1.14 | 8.5 |
| ospfHostCfgAreaID | 1.3.6.1.2.1.14.6.1.6 | 8.5 |
| ospfIfLsaCount | 1.3.6.1.2.1.14.7.1.21 | 8.5 |
| ospfIfLsaCksumSum | 1.3.6.1.2.1.14.7.1.22 | 8.5 |
| ospfIfDesignatedRouterId | 1.3.6.1.2.1.14.7.1.23 | 8.5 |
| ospfIfBackupDesignatedRouterId | 1.3.6.1.2.1.14.7.1.24 | 8.5 |
| ospfVirtIfLsaCount | 1.3.6.1.2.1.14.9.1.12 | 8.5 |
| ospfVirtIfLsaCksumSum | 1.3.6.1.2.1.14.9.1.13 | 8.5 |
| ospfNbrRestartHelperStatus | 1.3.6.1.2.1.14.10.1.12 | 8.5 |
| ospfVirtNbrRestartHelperAge | 1.3.6.1.2.1.14.11.1.10 | 8.5 |
| ospfVirtNbrRestartHelperExitReason | 1.3.6.1.2.1.14.11.1.11 | 8.5 |
| ospfAreaAggregateExtRouteTag | 1.3.6.1.2.1.14.14.1.7 | 8.5 |
| ospfLocalLsdbTable | 1.3.6.1.2.1.14.17 | 8.5 |
| ospfLocalLsdbEntry | 1.3.6.1.2.1.14.17.1 | 8.5 |
| ospfLocalLsdbIpAddress | 1.3.6.1.2.1.14.17.1.1 | 8.5 |
| ospfLocalLsdbAddressLessIf | 1.3.6.1.2.1.14.17.1.2 | 8.5 |
| ospfLocalLsdbType | 1.3.6.1.2.1.14.17.1.3 | 8.5 |
| ospfLocalLsdbLsid | 1.3.6.1.2.1.14.17.1.4 | 8.5 |
| ospfLocalLsdbRouterId | 1.3.6.1.2.1.14.17.1.5 | 8.5 |
| ospfLocalLsdbSequence | 1.3.6.1.2.1.14.17.1.6 | 8.5 |
| ospfLocalLsdbAge | 1.3.6.1.2.1.14.17.1.7 | 8.5 |
| ospfLocalLsdbChecksum | 1.3.6.1.2.1.14.17.1.8 | 8.5 |
| ospfLocalLsdbAdvertisement | 1.3.6.1.2.1.14.17.1.9 | 8.5 |
| ospfVirtLocalLsdbTable | 1.3.6.1.2.1.14.18 | 8.5 |
| ospfVirtLocalLsdbEntry | 1.3.6.1.2.1.14.18.1 | 8.5 |
| ospfVirtLocalLsdbTransitArea | 1.3.6.1.2.1.14.18.1.1 | 8.5 |
| ospfVirtLocalLsdbNeighbor | 1.3.6.1.2.1.14.18.1.2 | 8.5 |
| ospfVirtLocalLsdbType | 1.3.6.1.2.1.14.18.1.3 | 8.5 |
| ospfVirtLocalLsdbLsid | 1.3.6.1.2.1.14.18.1.4 | 8.5 |

Table 32: Common (continued)

| Object Name | Object OID | New in VOSS Release |
| --- | --- | --- |
| ospfVirtLocalLsdbRouterId | 1.3.6.1.2.1.14.18.1.5 | 8.5 |
| ospfVirtLocalLsdbSequence | 1.3.6.1.2.1.14.18.1.6 | 8.5 |
| ospfVirtLocalLsdbAge | 1.3.6.1.2.1.14.18.1.7 | 8.5 |
| ospfVirtLocalLsdbChecksum | 1.3.6.1.2.1.14.18.1.8 | 8.5 |
| ospfVirtLocalLsdbAdvertisement | 1.3.6.1.2.1.14.18.1.9 | 8.5 |
| ospfAsLsdbTable | 1.3.6.1.2.1.14.19 | 8.5 |
| ospfAsLsdbEntry | 1.3.6.1.2.1.14.19.1 | 8.5 |
| ospfAsLsdbType | 1.3.6.1.2.1.14.19.1.1 | 8.5 |
| ospfAsLsdbLsid | 1.3.6.1.2.1.14.19.1.2 | 8.5 |
| ospfAsLsdbRouterId | 1.3.6.1.2.1.14.19.1.3 | 8.5 |
| ospfAsLsdbSequence | 1.3.6.1.2.1.14.19.1.4 | 8.5 |
| ospfAsLsdbAge | 1.3.6.1.2.1.14.19.1.5 | 8.5 |
| ospfAsLsdbChecksum | 1.3.6.1.2.1.14.19.1.6 | 8.5 |
| ospfAsLsdbAdvertisement | 1.3.6.1.2.1.14.19.1.7 | 8.5 |
| ospfAreaLsaCountTable | 1.3.6.1.2.1.14.20 | 8.5 |
| ospfAreaLsaCountEntry | 1.3.6.1.2.1.14.20.1 | 8.5 |
| ospfAreaLsaCountAreaId | 1.3.6.1.2.1.14.20.1.1 | 8.5 |
| ospfAreaLsaCountLsaType | 1.3.6.1.2.1.14.20.1.2 | 8.5 |
| ospfAreaLsaCountNumber | 1.3.6.1.2.1.14.20.1.3 | 8.5 |
| rcIpPimGlobalSptInfiniteThreshold | 1.3.6.1.4.1.2272.1.8.16.1.13 | 8.5 |
| rcIpv6PimGlobalSptInfiniteThreshold | 1.3.6.1.4.1.2272.1.8.112.1.1.12 | 8.5 |
| rcIsisGlobalIpTunnelOverlay | 1.3.6.1.4.1.2272.1.63.1.34 | 8.6 |
| rcVrfActive | 1.3.6.1.4.1.2272.1.203.1.1.1.2.1.13 | 8.6 |

Table 33: 5520 Series

| Object Name | Object OID | New in Release |
| --- | --- | --- |
| rcAutoSense | 1.3.6.1.4.1.2272.1.231 | 8.5 |
| rcAutoSenseMib | 1.3.6.1.4.1.2272.1.231.1 | 8.5 |
| rcAutoSenseNotifications | 1.3.6.1.4.1.2272.1.231.1.0 | 8.5 |
| rcAutoSenseObjects | 1.3.6.1.4.1.2272.1.231.1.1 | 8.5 |
| rcAutoSenseScalars | 1.3.6.1.4.1.2272.1.231.1.1.1 | 8.5 |
| rcAutoSenseAccessDiffservEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.1 | 8.5 |
| rcAutoSenseDataIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.2 | 8.5 |

**Table 33: 5520 Series (continued)**

| Object Name | Object OID | New in Release |
|---|---|---|
| rcAutoSenseEapolVoiceLldpAuthEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.3 | 8.5 |
| rcAutoSenseFaMsgAuthEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.4 | 8.5 |
| rcAutoSenseFaAuthenticationKey | 1.3.6.1.4.1.2272.1.231.1.1.1.5 | 8.5 |
| rcAutoSenseIsisHelloAuthType | 1.3.6.1.4.1.2272.1.231.1.1.1.6 | 8.5 |
| rcAutoSenseIsisHelloAuthKey | 1.3.6.1.4.1.2272.1.231.1.1.1.7 | 8.5 |
| rcAutoSenseOnboardingIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.8 | 8.5 |
| rcAutoSenseQos8021pOverrideEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.9 | 8.5 |
| rcAutoSenseVoiceIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.10 | 8.5 |
| rcAutoSenseVoiceCvid | 1.3.6.1.4.1.2272.1.231.1.1.1.11 | 8.5 |
| rcAutoSenseIsisHelloAuthKeyId | 1.3.6.1.4.1.2272.1.231.1.1.1.12 | 8.5 |
| rcAutoSenseDhcpDetection | 1.3.6.1.4.1.2272.1.231.1.1.1.13 | 8.5 |
| rcAutoSenseFaCameraIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.14 | 8.5 |
| rcAutoSenseFaProxyMgmtIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.15 | 8.5 |
| rcAutoSenseFaProxyMgmtCvid | 1.3.6.1.4.1.2272.1.231.1.1.1.16 | 8.5 |
| rcAutoSenseFaProxyNoAuthIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.17 | 8.5 |
| rcAutoSenseFaVirtualSwitchIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.18 | 8.5 |
| rcAutoSenseFaWapType1Isid | 1.3.6.1.4.1.2272.1.231.1.1.1.19 | 8.5 |
| rcAutoSenseFaCameraEapolStatus | 1.3.6.1.4.1.2272.1.231.1.1.1.20 | 8.5 |
| rcAutoSenseFaEapolOVSStatus | 1.3.6.1.4.1.2272.1.231.1.1.1.21 | 8.5 |
| rcAutoSenseFaEapolWap1Status | 1.3.6.1.4.1.2272.1.231.1.1.1.22 | 8.5 |
| rcAutoSenseWaitInterval | 1.3.6.1.4.1.2272.1.231.1.1.1.23 | 8.5 |
| rcPortAutoSenseDataIsid | 1.3.6.1.4.1.2272.1.4.10.1.1.136 | 8.5 |
| rcNlsMgmtInterfaceDropIcmpFragEnable | 1.3.6.1.4.1.2272.1.223.1.1.16 | 8.5 |
| rcNlsMgmtInterfaceDropIcmpv6FragEnable | 1.3.6.1.4.1.2272.1.223.1.1.17 | 8.5 |
| rcNlsMgmtIsid | 1.3.6.1.4.1.2272.1.223.1.1.18 | 8.5 |
| rcIpConfGlobalTcpAdjustMssEnable | 1.3.6.1.4.1.2272.1.8.1.6.29 | 8.5 |
| rcIpConfGlobalTcpAdjustMssStatus | 1.3.6.1.4.1.2272.1.8.1.6.30 | 8.5 |

Table 33: 5520 Series (continued)

| Object Name | Object OID | New in Release |
|---|---|---|
| rcIpConfGlobalTcpAdjustMssType | 1.3.6.1.4.1.2272.1.8.1.6.31 | 8.5 |
| rcIpConfGlobalTcpAdjustMssValue | 1.3.6.1.4.1.2272.1.8.1.6.32 | 8.5 |

Table 34: 5420 Series

| Object Name | Object OID | New in Release |
|---|---|---|
| rcAutoSense | 1.3.6.1.4.1.2272.1.231 | 8.5 |
| rcAutoSenseMib | 1.3.6.1.4.1.2272.1.231.1 | 8.5 |
| rcAutoSenseNotifications | 1.3.6.1.4.1.2272.1.231.1.0 | 8.5 |
| rcAutoSenseObjects | 1.3.6.1.4.1.2272.1.231.1.1 | 8.5 |
| rcAutoSenseScalars | 1.3.6.1.4.1.2272.1.231.1.1.1 | 8.5 |
| rcAutoSenseAccessDiffservEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.1 | 8.5 |
| rcAutoSenseDataIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.2 | 8.5 |
| rcAutoSenseEapolVoiceLldpAuthEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.3 | 8.5 |
| rcAutoSenseFaMsgAuthEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.4 | 8.5 |
| rcAutoSenseFaAuthenticationKey | 1.3.6.1.4.1.2272.1.231.1.1.1.5 | 8.5 |
| rcAutoSenseIsisHelloAuthType | 1.3.6.1.4.1.2272.1.231.1.1.1.6 | 8.5 |
| rcAutoSenseIsisHelloAuthKey | 1.3.6.1.4.1.2272.1.231.1.1.1.7 | 8.5 |
| rcAutoSenseOnboardingIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.8 | 8.5 |
| rcAutoSenseQos8021pOverrideEnable | 1.3.6.1.4.1.2272.1.231.1.1.1.9 | 8.5 |
| rcAutoSenseVoiceIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.10 | 8.5 |
| rcAutoSenseVoiceCvid | 1.3.6.1.4.1.2272.1.231.1.1.1.11 | 8.5 |
| rcAutoSenseIsisHelloAuthKeyId | 1.3.6.1.4.1.2272.1.231.1.1.1.12 | 8.5 |
| rcAutoSenseDhcpDetection | 1.3.6.1.4.1.2272.1.231.1.1.1.13 | 8.5 |
| rcAutoSenseFaCameraIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.14 | 8.5 |
| rcAutoSenseFaProxyMgmtIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.15 | 8.5 |
| rcAutoSenseFaProxyMgmtCvid | 1.3.6.1.4.1.2272.1.231.1.1.1.16 | 8.5 |
| rcAutoSenseFaProxyNoAuthIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.17 | 8.5 |
| rcAutoSenseFaVirtualSwitchIsid | 1.3.6.1.4.1.2272.1.231.1.1.1.18 | 8.5 |
| rcAutoSenseFaWapType1Isid | 1.3.6.1.4.1.2272.1.231.1.1.1.19 | 8.5 |
| rcAutoSenseFaCameraEapolStatus | 1.3.6.1.4.1.2272.1.231.1.1.1.20 | 8.5 |
| rcAutoSenseFaEapolOVSStatus | 1.3.6.1.4.1.2272.1.231.1.1.1.21 | 8.5 |
| rcAutoSenseFaEapolWap1Status | 1.3.6.1.4.1.2272.1.231.1.1.1.22 | 8.5 |
| rcAutoSenseWaitInterval | 1.3.6.1.4.1.2272.1.231.1.1.1.23 | 8.5 |

Table 34: 5420 Series (continued)

| Object Name | Object OID | New in Release |
|---|---|---|
| rcPortAutoSenseDataIsid | 1.3.6.1.4.1.2272.1.4.10.1.1.136 | 8.5 |
| rcNlsMgmtInterfaceDropIcmpFragEnable | 1.3.6.1.4.1.2272.1.223.1.1.16 | 8.5 |
| rcNlsMgmtInterfaceDropIcmpv6FragEnable | 1.3.6.1.4.1.2272.1.223.1.1.17 | 8.5 |
| rcNlsMgmtIsid | 1.3.6.1.4.1.2272.1.223.1.1.18 | 8.5 |
| rcIpConfGlobalTcpAdjustMssEnable | 1.3.6.1.4.1.2272.1.8.1.6.29 | 8.5 |
| rcIpConfGlobalTcpAdjustMssStatus | 1.3.6.1.4.1.2272.1.8.1.6.30 | 8.5 |
| rcIpConfGlobalTcpAdjustMssType | 1.3.6.1.4.1.2272.1.8.1.6.31 | 8.5 |
| rcIpConfGlobalTcpAdjustMssValue | 1.3.6.1.4.1.2272.1.8.1.6.32 | 8.5 |
| rc2kBootConfigEnableSpbmNodeScaling | 1.3.6.1.4.1.2272.1.100.5.1.63 | 8.6 |

Table 35: 5320 Series

| Object Name | Object OID | New in Release |
|---|---|---|
| rc2kBootConfigEnableSpbmNodeScaling | 1.3.6.1.4.1.2272.1.100.5.1.63 | 8.6 |