Switch Engine™ v34.1.1 Release Notes

New Features, Improvements, and Known Issues

9041281-00 Rev AA
November 2026



Copyright © 2026 All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: https://www.extremenetworks.com/support/policies/open-source-declaration/

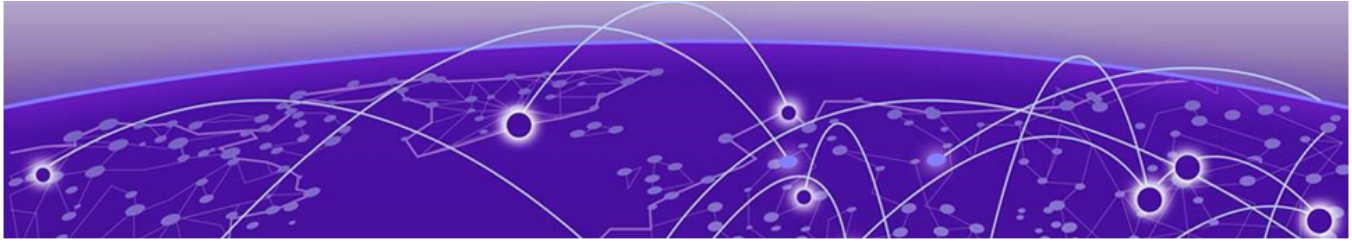


Table of Contents

Abstract vi
Preface vii
Text Conventions vii
Documentation and Training viii
Open Source Declarations ix
Training ix
Help and Support ix
Subscribe to Product Announcements x
Extreme Support Portal x
Send Feedback xii

Overview 13

Security Information 14
Linux Kernel 14
OpenSSL Version 14

Upgrading Switch Engine 15

Newly Purchased Switches Require Software Upgrade 16

Default Settings 17

Switch Engine Image File Names 20

New and Corrected Features 21
AVB Support on 5320-16P-2MXT-2X 21
Debug Data Local Staging and SCP Transfer Support 22
Modified CLI Commands 22
Supported Platforms 22
Notes 23
EAPS Flush Method Configuration 23
Configuration Syntax 23
Supported Platforms 24
Improved OSPF Scale on 5520, 5720, 7520, and 7720 Series Switches 24
OSPF Scale Limits 24
Supported Platforms 24
IPv6 First-Hop Security 24
New CLI Commands 25
Monitoring 26
Limitation 27
Supported Platforms 27
IPv6 Prefixes Exceeding 64 Bits 27
Modified CLI Commands 28
Supported Platforms 28

Fabric Attach Client Auto-Enable Authentication on Uplinks 28
Monitoring 28
Supported Platforms 28
Fabric Attach Clients That Fail Authentication Visibility 29
Monitoring 29
Supported Platforms 29
Fabric Attach Zero-Touch Client: Bypass NetLogin Authentication 29
Modified CLI Commands 29
Monitoring 30
Supported Platforms 30
Fan Speed Display on Silent-Mode Platforms 30
Modified CLI Commands 30
Supported Platforms 31
FEC Statistics Detailed Display 31
Modified CLI Commands 31
Supported Platforms 31
FIPS Compliance: Microsoft Golang for OpenSSL Support 31
Affected Platforms 31
Flood Rate-Limit Bandwidth Mode 32
Modified CLI Commands 32
Monitoring 32
Supported Platforms 32
Hybrid PQC Support in RADIUS TLS and Syslog TLS 32
Verification 33
Supported Platforms 33
MAC Locking Last Violating MAC Address Clear 33
New CLI Commands 33
Supported Platforms 33
Network Zones in ACL Policy Rules 34
Modified CLI Commands 34
Supported Platforms 36
OpenLDAP Upgrade to 2.6.12 36
Supported Platforms 36
NetLogin Client IP Tracking for RADIUS Framed-IP-Address 36
RADIUS Accounting Behavior 37
Configuration 37
Supported Platforms 37
Power Supply Failure Alert Severity Change 37
Affected Platforms 37
PQC Algorithms in OpenSSH 37
New CLI Commands 38
Monitoring 38
Supported Platforms 38
SNMP Interface Pause Frame Counters 38
New CLI Commands 39
Supported Platforms 39
SNMPv3 User Authentication with SHA-2 39
Modified CLI Commands 39
Migration from SHA-1 40

Supported Platforms....40
Transparent Clock Support on 5120 Series....40
CLI Commands....40
Monitoring....40
Supported Platforms....41
Update Time-Zone Information for UTC Changes in Alberta and British Columbia....41
Configuration....41
Supported Platforms....41
Updates to Avoid NTP Epocalypse - Y2036 Overflow....41
Supported Platforms....41
Updates to Avoid Epocalypse - Y2038 Overflow....42
Supported Platforms....42

Changing the Network Operating System....43
Making Your Initial Network Operating System Selection....43
Changing Your Network Operating System....44

ExtremeCloud IQ Agent Support....45

Extreme Hardware/Software Compatibility and Recommendation Matrices....48

Compatibility with Extreme Management Center....49

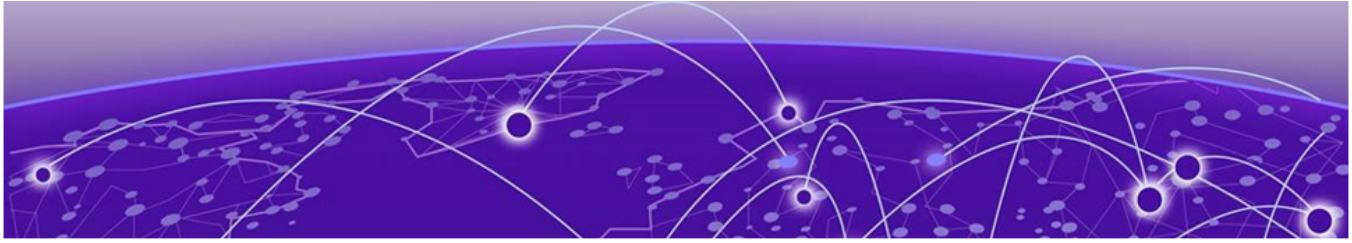
Supported MIBs....50

Tested Third-Party Products....51
Tested RADIUS Servers....51

Extreme Switch Security Assessment....52
DoS Attack Assessment....52
ICMP Attack Assessment....52
Port Scan Assessment....52

Limits....53
Limits Overview....53
Base License Limits....56
Premier License Limits....93
Notes for Limits Tables....103

Open Issues, Known Behaviors, and Resolved Issues....105
Open Issues....105
Known Behaviors....106
Resolved Issues in Switch Engine....107



Abstract

The release notes for Switch Engine™ v34.1.1 detail new features, software improvements, supported platforms, upgrade prerequisites, scaling limits, and known issues for Universal Hardware switches. New security features include IPv6 First-Hop Security with RA Guard, DHCPv6 snooping, IPv6 Source Guard, and Neighbor Discovery Inspection; hybrid post-quantum key exchange for SSH, RADIUS over TLS, and Syslog over TLS; and SNMPv3 authentication with SHA-256, SHA-384, and SHA-512. Access and policy enhancements include ACL rules that match both source and destination network zones, Fabric Attach bypass of NetLogin for zero-touch clients, visibility of Fabric Attach clients that fail authentication, automatic Fabric Attach authentication on uplinks, and NetLogin tracking of client IP addresses for the RADIUS Framed-IP-Address attribute. Other additions include a bandwidth mode for flood rate limiting, a configurable EAPS flush method, local staging of debug data for SCP transfer, detailed FEC statistics, SNMP pause frame counters, Transparent Clock support on the 5120 Series, improved OSPF scale on the 5520, 5720, 7520, and 7720 Series, AVB support on the 5320-16P-2MXT-2X, and updates for the Y2038 and Y2036 time overflows and for time-zone changes in Alberta, the Northwest Territories, and British Columbia. Infrastructure upgrades include Linux kernel 6.12 and OpenSSL 3.5.8. Supported deployment models include dual-OS capability with automatic selection by using ExtremeCloud IQ, Extreme Management Center, or manual bootloader configuration. Base and Premier license limits, compatibility matrices with Extreme Management Center, supported RADIUS servers, and tested third-party products are specified. Open issues, known behavioral limitations such as DHCPv6 snooping binding learning when only IPv4 source IP lockdown is enabled, and resolved defects are documented for technical implementation and troubleshooting."

Preface

Read the following topics to learn about:

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

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as Extreme Networks switches, the product is referred to as *the switch*.

Table 1: Notes and warnings

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

Table 2: Text

Convention	Description
`screen displays`	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words *enter* and *type*	When you see the word *enter* in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says *type*.
Key names	Key names are written in boldface, for example **Ctrl** or **Esc**. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press **Ctrl**+**Alt**+**Del**
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
italic text	Italic text indicates variable content.
`[ ]`	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
`{ x | y | z }`	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation
Release Notes

Hardware and Software Compatibility for Extreme Networks products
Extreme Optics Compatibility
Other Resources such as articles, white papers, and case studies

Open Source Declarations

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

Training

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

Help and Support

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

Extreme Agent ONE™ Coworker "Talk to Support"

For customers with Extreme Platform ONE™ ExtremeWorks / ExtremeWorks Premier subscriptions, technical support cases can be created using Extreme Agent ONE™ Coworker. This option requires that a user has an existing Extreme Networks Support Portal account login (using the same email address as used to log into Extreme Platform ONE™ Networking.) For more information about Extreme Agent ONE™ Coworker “Talk to Support”, see the Extreme Agent ONE™ Coworker documentation.

Extreme Portal

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

The Hub

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

Call GTAC

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

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

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)

- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

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

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

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

Extreme Support Portal

The Extreme Networks Extreme Support Portal provides access to:

- License management
- Firmware downloads
- Case management
- GTAC Knowledge base
- And more

To access the portal, you must first create an account. The Extreme Support Portal provides separate registration forms for End Customers and Partners.

Before you begin, gather the following: A valid corporate email address — required to receive the appropriate level of portal access. Your company details: company name, address, phone number (include country code if outside the US), job title, and job role. At least one supporting identifier: Sold-to number, serial number, contract number, or Purchase Order (PO) number.

Create an End Customer Extreme Support Portal Account

The Extreme Support Portal provides separate registration forms for End Customers and Partners. Use this task to create an End Customer account.

Note
If you are adding a new user to an existing account, the new user must complete and be vetted through the same registration process.

The portal registration page is at https://extreme-networks.my.site.com/.

1. Go to the Extreme Support Portal and select **Login or Signup** on the login page.
2. On the login page, select **Create a new Extreme Customer Portal Account** .

3. Complete all required fields (marked *). Use your corporate email address to ensure the appropriate access level.

   Important
   Do not register with a personal email address.

4. Read the Support Portal Terms of Use and select the acknowledgment check box to accept.

   Important
   You must accept the Terms of Use before you can submit the form. You can adjust or withdraw consent at any time through the Email Preference Center.

5. Select **Register**.

   A confirmation message appears on the page and a password setup email is sent to your corporate email address.

6. Open the password setup email and follow the instructions to create your portal password.

   Tip
   If the email does not appear in your inbox within a few minutes, check your spam or junk mail folder.

7. To log in to the support portal go to: https://extremesaas.my.site.com/.

Create a Partner Extreme Support Portal Account

The Extreme Support Portal provides separate registration forms for End Customers and Partners. Use this task to create a Partner account.

The portal registration page is at https://extreme-networks.my.site.com/.

1. Go to the Extreme Support Portal and select **Login or Signup** on the login page.
2. On the login page, select **Create a new Extreme Partner Portal Account**.
3. Complete all required fields (marked *). Use your corporate email address to ensure the appropriate access level.

   Important
   Do not register with a personal email address.

4. Read the following terms and select the required check boxes:
   - Support Portal Terms of Use
   - Extreme Networks Standard Reseller Terms and Conditions
   - Trade Compliance Terms and Conditions

   You can adjust or withdraw consent at any time through the Email Preference Center.

5. Select **Register**.

   A confirmation message appears on the page and a password setup email is sent to your corporate email address.

6. Open the password setup email and follow the instructions to create your portal password.



Tip
If the email does not appear in your inbox within a few minutes, check your spam or junk mail folder.

7. To log in to the support portal go to: https://extremesaas.my.site.com/.

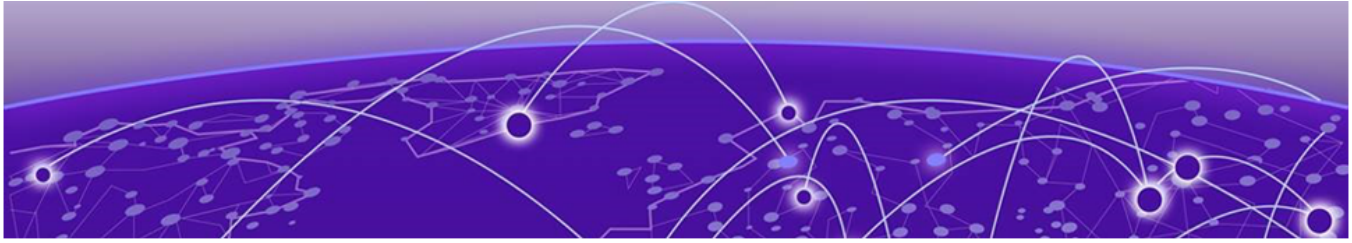
Send Feedback

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

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

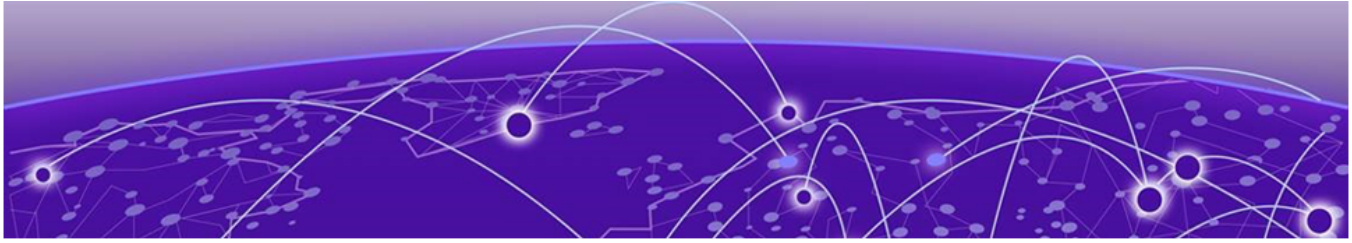
To send feedback, email us at Product-Documentation@extremenetworks.com.

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



Overview

These release notes document the title version of Switch Engine, which adds features and resolves software deficiencies.



Security Information

Linux Kernel on page 14
OpenSSL Version on page 14

The following section covers important security information.

Linux Kernel

Switch Engine 34.1.1 uses the following Linux Kernel versions by platform:

- **Linux Kernel 6.12:** 4220, 5320, 5420, 5520, 5720, 7520, 7720, Virtual Machine, Digital Twin
- **Linux Kernel 5.10:** 4120, 5120

OpenSSL Version

This version of Switch Engine uses OpenSSL FIPS Object Module 3.5.8.

Upgrading Switch Engine

For instructions about upgrading software, see *Software Upgrade and Boot Options* in the user guide.

A Switch Engine core image (.xos file) must be downloaded and installed on the alternate (non-active) partition. If you try to download to an active partition, the system displays the following error message: `Error: Image can only be installed to the non-active partition.`. A modular software package (.xmod file) can still be downloaded and installed on either the active or alternate partition.



Note

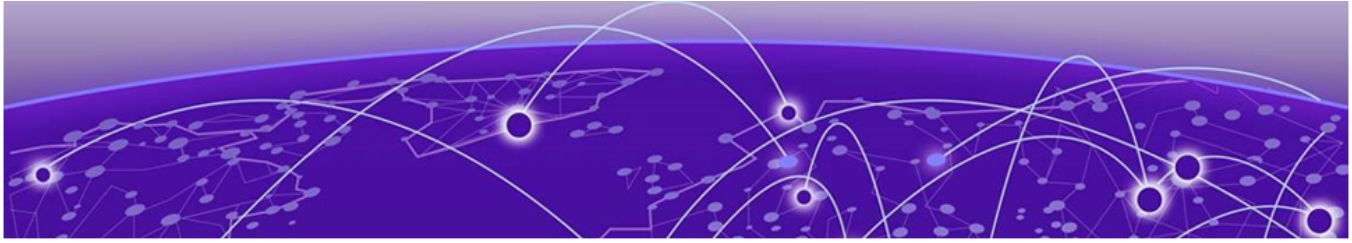
New 5420 and 5520 PoE switches use a new version of the PoE microcontroller that prevents the switch from downgrading to older versions and prevents operating system switchover to unsupported VOSS versions.

The following error message is displayed during the downgrades to older versions:

```
Error: Failed to download image - summit_arm-31.6.1.3.xos does not
include compatible PoE microcontroller support. See the User Guide for
information on installing a newer software release. See the
Hardware/Software Compatibility and Recommendation Matrices to verify the
supported releases.
```

5420 and 5520 PoE switches that use a new version of the PoE microcontroller can be identified for by checking the PoE firmware revision (5.0 or later) by entering the `show inline-power stats` command (line four):

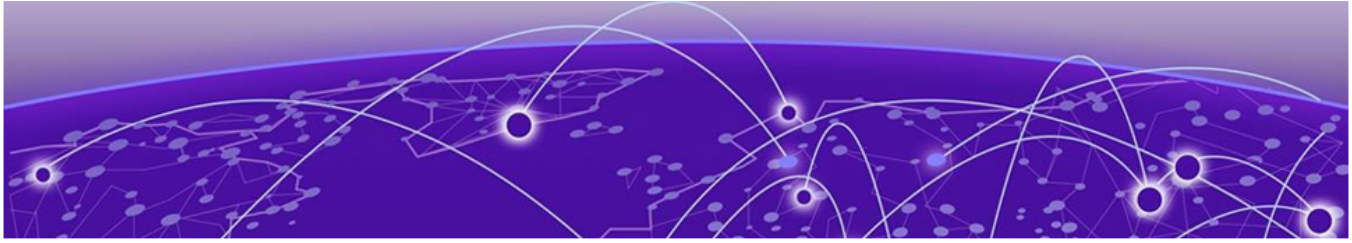
```
# show inline-power stats
Inline-Power Slot Statistics
Firmware status               : Operational
Firmware revision             : 5.0.0b4
Total ports powered           : 3
Total ports awaiting power    : 20
Total ports faulted           : 0
Total ports disabled          : 1
```



Newly Purchased Switches Require Software Upgrade

Newly delivered switches typically have pre-GA (general availability) software installed. You should promptly upgrade the software to the latest version available by visiting the Extreme Portal.

For information about upgrading the software, see the *Switch Engine Upgrade Process* topic in the *Software Upgrade and Boot Options* chapter of the user guide.



Default Settings

The following table shows the default settings for Switch Engine starting with version 31.6, and shows any changes that have been made to these settings and in what version these changes were made.

Table 4: Default Settings

Feature	31.6 and later	32.4 and later
1G behavior in 10G ports (5420 and 5520 series switches)	Autoneg OFF for port when 1G optic is inserted in a 10G port	
Account Lockout	After 3 consecutive login failures, account is locked for 5 minutes. [a]	
Auto-Discovery for Universal Hardware	Enabled.	
AVB	Disabled.	
BFD Strict Session Protection	Disabled.	
BGP	Disabled.	
Bluetooth	Enabled.	
BOOTP Relay	Disabled.	
CDP	Enabled.	
Configuration auto save	Disabled.	
Clear-flow	Disabled.	
Diagnostics	Admin level privileges required to show diagnostics. [a]	
DHCP	Disabled.	
DNS Cache Resolver and Analytics	Disabled.	
IPFIX	Disabled.	
IP NAT	Disabled.	
EAPS	Disabled.	
EDP	Enabled.	
ELRP	Disabled.	
ESRP	Disabled.	

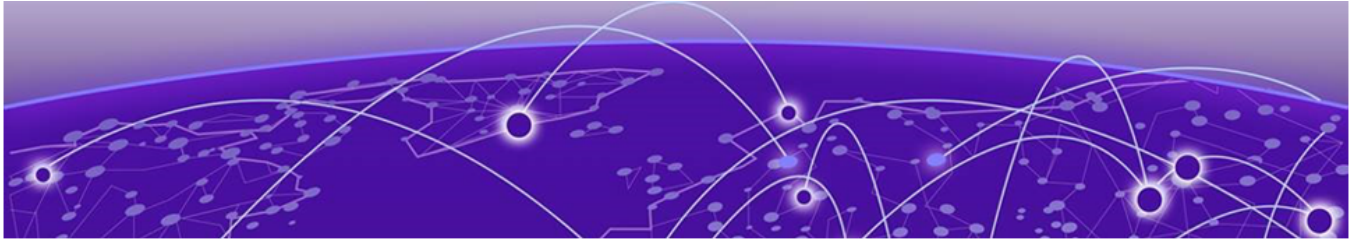
[a] If you choose enhanced security mode when initially setting up the switch or after running `unconfigure switch all`.

Table 4: Default Settings (continued)

Feature	31.6 and later	32.4 and later
Extended Edge Switching (VPEX)	Disabled.	
ExtremeCloud IQ	Enabled	
FEC	Enabled on Native 25Gb ports.	
Identity Management	Disabled.	
IGMP	Enabled, set to IGMPv2 compatibility mode.	
IGMP Snooping	Enabled.	
Image Integrity Check	Disabled.	
IP Route Compression	Enabled.	
ISIS	Disabled.	
LLDP	Enabled.	
Log	Admin level privileges required to show log. [a]	
Logging memory buffer	Generate an event when the logging memory buffer exceeds 90% of capacity. [a]	
MAC Security	Disabled.	
MLD	Disabled.	
MLD Snooping	Disabled.	
MPLS	Disabled.	
MSRP	Disabled.	
MSTP	Enabled.	
NetLogin	All types of authentication are disabled.	
NTP	Disabled.	
ONEPolicy	Disabled.	
Policy rule model	Hierarchical (Unless upgrading from 30.5 with a saved configuration set to access list.)	
OpenFlow	Disabled.	
OSPF	Disabled.	
OVSDB	Disabled.	
Passwords	Plain text password entry not allowed. [a]	
PIM	Disabled.	
PIM Snooping	Disabled.	

Table 4: Default Settings (continued)

Feature	31.6 and later	32.4 and later
PoE Fast PoE Perpetual PoE	Enabled. Disabled. Disabled.	
RADIUS	Disabled for both switch management and network login.	
RIP	Disabled.	
RMON	Disabled. However, even in the disabled state, the switch responds to RMON queries and sets for alarms and events.	
sFlow	Disabled.	
SNMP server	Disabled. [a]	
SSH	Disabled.	
Stacking-support	Enabled.	Disabled for 5120, Extreme 7520, and 7720 only.
Stacking auto-discovery	Enabled.	
STP	Enabled.	
Syslog	Disabled.	
TACACS	Disabled.	
Telnet	Enabled. [a]	
VPEX IP Multicast Replication	BPE	
VPLS	All newly created VPLS instances are enabled.	
Watchdog	Enabled.	
Web HTTP server	Enabled. [a]	
Web HTTPS server	Enabled. [a]	

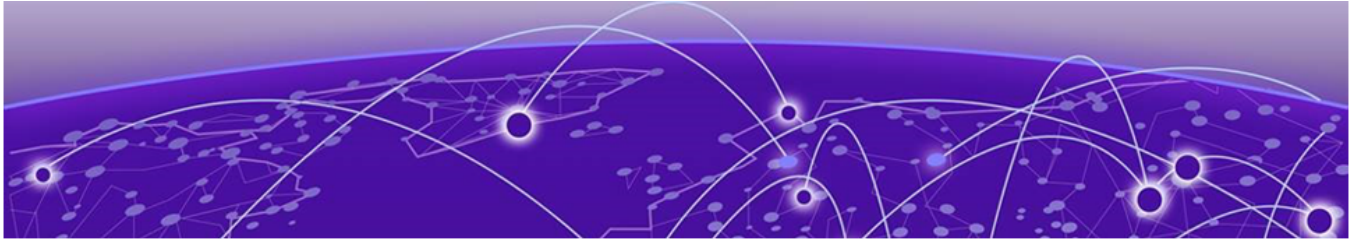


Switch Engine Image File Names

You can identify the appropriate image or module for your platform based on the file name prefix of the image.

Table 5: Switch Engine Image Types (Prefixes)

Switches	Image File Type (Prefix)
4120, 5120	`rzg2` Example: `rzg2-33.3.1.x.xos`
4220, 5320, 5420, 5520	`summit_arm` Example: `summit_arm-33.1.1.x.xos`
5720, 7520, 7720	`onie` Example: `onie-33.1.1.6.x86_64.xos`



New and Corrected Features

AVB Support on 5320-16P-2MXT-2X on page 21
Debug Data Local Staging and SCP Transfer Support on page 22
EAPS Flush Method Configuration on page 23
Improved OSPF Scale on 5520, 5720, 7520, and 7720 Series Switches on page 24
IPv6 First-Hop Security on page 24
IPv6 Prefixes Exceeding 64 Bits on page 27
Fabric Attach Client Auto-Enable Authentication on Uplinks on page 28
Fabric Attach Clients That Fail Authentication Visibility on page 29
Fabric Attach Zero-Touch Client: Bypass NetLogin Authentication on page 29
Fan Speed Display on Silent-Mode Platforms on page 30
FEC Statistics Detailed Display on page 31
FIPS Compliance: Microsoft Golang for OpenSSL Support on page 31
Flood Rate-Limit Bandwidth Mode on page 32
Hybrid PQC Support in RADIUS TLS and Syslog TLS on page 32
MAC Locking Last Violating MAC Address Clear on page 33
Network Zones in ACL Policy Rules on page 34
OpenLDAP Upgrade to 2.6.12 on page 36
NetLogin Client IP Tracking for RADIUS Framed-IP-Address on page 36
Power Supply Failure Alert Severity Change on page 37
PQC Algorithms in OpenSSH on page 37
SNMP Interface Pause Frame Counters on page 38
SNMPv3 User Authentication with SHA-2 on page 39
Transparent Clock Support on 5120 Series on page 40
Update Time-Zone Information for UTC Changes in Alberta and British Columbia on page 41
Updates to Avoid NTP Epocalypse - Y2036 Overflow on page 41
Updates to Avoid Epocalypse - Y2038 Overflow on page 42

This section lists the new and corrected features supported in this version:

AVB Support on 5320-16P-2MXT-2X

Version 34.1.1 enables Audio Video Bridging (AVB) functionality on the 5320-16P-2MXT-2X switch.

AVB functionality is included under the Base license and does not require additional licensing or configuration changes.

Debug Data Local Staging and SCP Transfer Support

Version 34.1.1 extends the upload debug command to support local staging of debug files for manual transfer via secure protocols such as SCP, providing an alternative to TFTP and enabling USB-based file transport.

Overview: The existing upload debug command generates and transfers debug tar files using TFTP in a single operation. Version 34.1.1 adds a stage-locally option that generates debug files and stages them on the switch internal storage (/usr/local/tmp) or on removable USB storage (/usr/local/ext), allowing users to manually transfer them later using SCP, TFTP, or by physically transporting USB media. This approach supports customers requiring secure communication and provides an alternative when direct network access is unavailable.

Debug files staged locally can be transferred manually, avoiding potential disk-full scenarios on the switch internal partition and enabling flexible offloading strategies.

Modified CLI Commands

Stage debug data locally for manual transfer:

upload debug stage-locally [*directory_path*|**default**]

When `default` is specified, debug files are staged in the directory configured by **configure debug core-dumps** *directory_path*, or in /usr/local/tmp if core-dumps is disabled.

Example — Stage to USB storage:

```
upload debug stage-locally /usr/local/ext
```

Optional filename formatting (applies to staged or uploaded files):

upload debug stage-locally *directory_path* **name-format** [**none**|**date**|**date-time**]

Name-format options: none (base filename only), date (append ISO 8601 date YYYY-mm-dd), date-time (append mmddHHMM, default).

Traditional TFTP upload (unchanged):

upload debug [*ipaddress*|*hostname*] {{**vr** *vrname*} {**block-size** *block_size*}}

Supported Platforms

All platforms.

Notes

- In a stack, the primary switch fetches and stores debug tar files from all Backup and Standby nodes, eliminating the need to connect to individual stack slots.
- After staging completes, a reminder is displayed to manually remove debug files using `rm` after transfer.

EAPS Flush Method Configuration

Ethernet Automatic Protection Switching (EAPS) now supports configurable flush methods to control which MAC addresses are cleared when a ring port fails. Select between port-based flushing (default) or VLAN-and-port-based flushing to optimize convergence timing and MAC relearning in multi-VLAN ring environments.

Issue: When an EAPS ring port fails, stale MAC address entries learned on that port must be cleared to force proper relearning on the alternate path. Previously, EAPS used a fixed flushing strategy that was suboptimal for multi-VLAN deployments, resulting in unnecessary convergence delays and delayed MAC relearning across VLANs not affected by the port failure.

Resolution: Version 34.1.1 introduces the `configure eaps flush-method` command with two configurable modes: `port-based` (flush only MAC addresses learned on the failed port) and `vlan-and-port-based` (flush only MAC addresses learned on the failed port within affected VLANs). This allows operators to reduce the learning scope during failover and improve convergence time in multi-VLAN ring topologies.

Key Capabilities:

- Configurable EAPS MAC flush method at the ring level using the `configure eaps flush-method` command
- Two flush modes: `port-based` (flush MACs on failed port) and `vlan-and-port-based` (flush MACs on failed port within affected VLANs)
- Port-based flushing is the default mode, maintaining backward compatibility with existing deployments
- VLAN-and-port-based mode reduces MAC clearing scope, improving convergence performance in multi-VLAN environments
- Monitoring commands `show eaps` and `show eaps detail` display the configured flush method
- Configuration persists across device reboot

Configuration Syntax

Configure the EAPS flush method:

```
configure eaps flush-method [port-based | vlan-and-port-based]
```

Example: Enable VLAN-and-Port-Based Flushing

```
configure eaps flush-method vlan-and-port-based
show eaps detail
```

This configuration optimizes convergence for multi-VLAN ring topologies by limiting MAC address clearing to only the failed port and affected VLANs.

Supported Platforms

All platforms.

Improved OSPF Scale on 5520, 5720, 7520, and 7720 Series Switches

Version 34.1.1 increases OSPF Link State Database (LSDB) capacity on the 5520, 5720, 7520, and 7720 Series switches, enabling support for larger-scale OSPF network deployments.

Overview: OSPF scale improvements on the 5520, 5720, 7520, and 7720 Series switches support significantly increased inter-area routes (Type-3 LSAs) and external routes (Type-5 LSAs). These enhancements enable customers migrating from legacy platforms to deploy OSPF networks with larger routing scales while maintaining stable operation.

OSPF Scale Limits

OSPF Inter-area Routes (Type-3):

- 5520 and 5720 Series: Up to 75,000 routes
- 7520 and 7720 Series: Up to 200,000 routes

OSPF External Routes (Type-5):

- 5520 and 5720 Series: Up to 75,000 routes
- 7520 and 7720 Series: Up to 200,000 routes

OSPF Inter-area Routes (Type-3) and External Routes (Type-5) Combined:

- 5520 and 5720 Series: Up to 150,000 routes
- 7520 and 7720 Series: Up to 250,000 routes

Supported Platforms

5520, 5720, 7520, and 7720 Series switches.

IPv6 First-Hop Security

Version 34.1.1 introduces IPv6 First-Hop Security features that protect edge ports from IPv6 attacks and unauthorized network access.

Overview: IPv6 First-Hop Security implements multiple edge protection mechanisms that work together to prevent IPv6-based attacks and enforce network access policies:

- **RA Guard:** Blocks rogue Router Advertisements on untrusted ports, preventing router impersonation and man-in-the-middle attacks.
- **DHCPv6 Snooping:** Blocks DHCPv6 server messages on untrusted ports, preventing rogue DHCPv6 servers. It also builds a binding table of IPv6 address, MAC address, VLAN and port from DHCPv6 exchanges, similar to DHCPv4 snooping.

- **IPv6 Source Guard:** Allows IPv6 traffic on a port only from source addresses in the binding table, learned through DHCPv6 snooping or configured as static bindings. Traffic from link-local source addresses is denied by default.
- **Neighbor Discovery Inspection (NDI):** Validates Neighbor Solicitation and Neighbor Advertisement messages against the binding table, preventing neighbor spoofing and neighbor cache poisoning attacks.

New CLI Commands

Enable RA Guard:

enable ip-security {**ipv6**} **ra-guard** {**vlan**} *vlan_name* **ports** [*port_list* | **all**] [**violation-action** [**drop-packet** {[**block-port** | **block-mac**] [**permanently** | **duration** *duration_in_seconds*]}]]

Disable RA Guard:

disable ip-security {**ipv6**} **ra-guard** {**vlan**} *vlan_name* **ports** [*port_list* | **all**]

Configure or unconfigure trusted ports for routers that send IPv6 Router Advertisements:

[**configure** | **unconfigure**] **trusted-ports** {**ipv6**} [*port_list* | **all**] **trust-for router**

Enable DHCPv6 snooping:

enable ip-security ipv6 dhcp-snooping {**vlan**} *vlan_name* **ports** [*port_list* | **all**] **violation-action** [**none** | **drop-packet** {[**block-port** | **block-mac**] [**permanently** | **duration** *duration_in_seconds*]}]

Disable DHCPv6 snooping:

disable ip-security ipv6 dhcp-snooping {**vlan**} *vlan_name* **ports** [*port_list* | **all**]

Clear DHCPv6 snooping entries:

clear ip-security ipv6 dhcp-snooping entries {**vlan**} *vlan_name* {**ipv6-address** [*ipv6_address*] | **ports** [*port_list*]}

Add a DHCPv6 binding entry:

configure ip-security dhcp-bindings add ipv6 *ipv6_address* **mac** *mac_address* {**vlan**} *vlan_name* **server-port** *server_port* **client-port** *client_port* **lease-time** *seconds*

Delete a DHCPv6 binding entry:

configure ip-security dhcp-bindings delete ipv6 *ipv6_address* {**vlan**} *vlan_name*

Configure or unconfigure trusted ports for DHCPv6 servers:

[configure | unconfigure] trusted-ports ipv6 [*port_list* | **all**] **trust-for dhcp-server**

Enable or disable DHCPv6 binding restoration:

[enable | disable] ip-security dhcp-bindings restoration {**ipv4-only** | **ipv6-only** | **both**}

Enable IPv6 Source Guard:

enable ip-security ipv6 source-ip-lockdown ports [*port_list* | **all**] {**link-local** [**permit** | **deny**]}

Disable IPv6 Source Guard:

disable ip-security ipv6 source-ip-lockdown ports [*port_list* | **all**]

Clear IPv6 Source Guard entries:

clear ip-security ipv6 source-ip-lockdown entries ports [*port_list* | **all**]

Enable Neighbor Discovery Inspection:

enable ip-security {**ipv6**} **nd-inspection source-mac destination-mac ip vlan** *vlan_name* **ports** [*port_list* | **all**] **violation-action drop-packet** {**block-port** [**permanently** | **duration** *duration_in_seconds*]}

Disable Neighbor Discovery Inspection:

disable ip-security {**ipv6**} **nd-inspection** {**vlan**} *vlan_name* **ports** [*port_list* | **all**]

Clear Neighbor Discovery Inspection violations:

clear ip-security {**ipv6**} **nd-inspection violations**

Monitoring

Show RA Guard configuration:

show ip-security {**ipv6**} **ra-guard** {**vlan**} *vlan_name*

Show RA Guard violations:

show ip-security {**ipv6**} **ra-guard violations** {**vlan**} *vlan_name*

Show DHCPv6 snooping configuration for a VLAN:

show ip-security ipv6 dhcp-snooping {**vlan**} *vlan_name*

Show DHCPv6 snooping entries for a VLAN:

show ip-security ipv6 dhcp-snooping entries {**vlan**} *vlan_name*

Show DHCPv6 snooping entries for a list of VLANs:

show ip-security ipv6 dhcp-snooping entries vlan *vlan_list*

Show DHCPv6 snooping violations for a VLAN:

show ip-security ipv6 dhcp-snooping violations {**vlan**} *vlan_name*

Show DHCPv6 snooping violations for a list of VLANs:

show ip-security ipv6 dhcp-snooping violations vlan *vlan_list*

Show IPv6 Source Guard:

show ip-security ipv6 source-ip-lockdown

Show Neighbor Discovery Inspection configuration:

show ip-security {**ipv6**} **nd-inspection** {**vlan**} *vlan_name*

Show Neighbor Discovery Inspection for a list of VLANs:

show ip-security {**ipv6**} **nd-inspection vlan** *vlan_list*

Show Neighbor Discovery Inspection violation count:

show ip-security {**ipv6**} **nd-inspection violations** {**vlan**} *vlan_name* {**ports** [*port_list* | **all**]}

Limitation

In a Layer 2 setup, if IPv4 source IP lockdown is enabled on a port and IPv6 source IP lockdown is not, DHCPv6 snooping does not learn bindings for the clients on that port. To avoid this, also enable IPv6 source IP lockdown on the same ports. Do not enable source IP lockdown on the port that faces the DHCPv6 server.

Supported Platforms

All platforms.

IPv6 Prefixes Exceeding 64 Bits

Version 34.1.1 enables support for IPv6 static routes and connected networks with prefix lengths exceeding 64 bits, accommodating modern IPv6 addressing schemes that require granular subnetting.

Overview: The switch now stores and forwards IPv6 routes with prefix lengths from /64 to /128 in the routing table, allowing more specific route matching and supporting point-to-point links with /127 or /128 prefixes.

Modified CLI Commands

Configure IPv6 static route with /127 prefix:

configure iproute add *vlan vlan* 2001:db8::/127 *nexthop_ip*

Configure IPv6 address on interface with /128 prefix:

configure *vlan vlan* **ipv6** address 2001:db8::1/128

View IPv6 routing table:

show iproute vr vr_name

Supported Platforms

All platforms with IPv6 routing support.

Fabric Attach Client Auto-Enable Authentication on Uplinks

Version 33.5.2 added automatic enablement of Fabric Attach message authentication on the uplinks of an FA Proxy switch towards the FA Server, if the FA Server is found to use Fabric Attach message authentication. Version 34.1.1 adds the same logic to automatically enable Fabric Attach message authentication on the uplinks of an FA Client switch towards the FA Proxy switch, if the FA Proxy is found to use Fabric Attach message authentication. This ensures that Fabric Attach with message authentication can be deployed even in cascaded deployments where two levels of FA switches (FA Proxy + FA Client) are connected behind the FA Server.

Overview: When Fabric Attach message authentication is enabled on the FA server ports, the FA Proxy switch connected into those ports will automatically enable Fabric Attach message authentication on those uplinks. When Fabric Attach message authentication is also enabled on all the other FA Proxy ports, an FA Client switch connected into any of those ports will automatically enable Fabric Attach message authentication on that uplink as well.

Monitoring

View Fabric Attach authentication status of ports:

show fabric attach ports *port_list* **authentication**

View detailed Fabric Attach authentication status of ports:

show fabric attach ports *port_list* **authentication detail**

Supported Platforms

All platforms with Fabric Attach support.

Fabric Attach Clients That Fail Authentication Visibility

Version 34.1.1 will now display FA Clients which either have no FA message authentication or fail FA message authentication on ports where FA message authentication is enabled. In previous software versions these clients were not shown in the FA elements table.

Overview: When Fabric Attach authentication is enabled on a port and the connected FA Client device does not have FA message authentication enabled (or it does but with the wrong authentication key) then none of the Fabric Attach information provided by the FA Client will be acted upon by the switch. The FA client will however now be shown in the FA elements table. The switch will always record the FA Client type and FA authentication outcome into the inbound RADIUS VSA FA-Client-Type (182) and VSA FA-Client-PSK (183) when sending the RADIUS authentication request for the client, so that the RADIUS server may still decide whether to allow the FA Client onto the network.

Supported Scenarios:

- Non-Extreme Fabric Attach clients connecting to Extreme Fabric Attach proxy switches
- RADIUS server returns Fabric Attach VSAs with VLAN and NSI mapping information

Monitoring

View Fabric Attach elements, including clients that have no FA message authentication or fail it:

```
show fabric attach elements
```

Supported Platforms

All platforms with Fabric Attach support.

Fabric Attach Zero-Touch Client: Bypass NetLogin Authentication

Version 34.1.1 adds an option to bypass NetLogin authentication for Fabric Attach Zero-Touch Client (ZTC) devices, allowing automatic network access provisioning without additional authentication steps.

Overview: When the bypass NetLogin option is enabled for a Fabric Attach ZTC client, the switch automatically provisions the client with network access through Fabric Attach without requiring the client to complete NetLogin authentication. This simplifies deployment for trusted devices in automated environments.

Modified CLI Commands

Enable NetLogin bypass for a Fabric Attach zero-touch client type:

```
configure fabric attach zero-touch-client client_type bypass-netlogin on
```

Monitoring

View the Fabric Attach zero-touch client configuration:

show fabric attach zero-touch-client

View NetLogin bypassed clients:

show fabric attach elements

show netlogin

show netlogin port *port_list*

Supported Platforms

All platforms with Fabric Attach support.

Fan Speed Display on Silent-Mode Platforms

Version 34.1.1 changes how the show fans command displays the fan speed on platforms that run their fans in silent mode. When the fan is off, the command now displays `Operational at 0 RPM`.

Overview: Some platforms run their fans in silent mode by design. The fan is intentionally off, or at a minimum baseline speed, when the system temperature is below a threshold. On these platforms, earlier releases displayed the baseline speed, for example 240 RPM. Beginning in 34.1.1, the fan state is still `Operational`, and the speed is displayed as 0 RPM when the fan is off. This display change accompanies the Linux kernel upgrade to version 6.12. `Operational at 0 RPM` is the standard display for silent mode.

Modified CLI Commands

Display the fan status:

show fans {**detail**}

Output in 34.1.1 on a silent-mode platform:

```
show fans
Slot-1 FanTray information:
State:                  Operational
NumFan:                 1
Airflow Direction:      Left to Right
Fan-1:                  Operational at 0 RPM
```

Output in earlier releases:

```
show fans
Slot-1 FanTray information:
State:                  Operational
NumFan:                 1
Airflow Direction:      Left to Right
Fan-1:                  Operational at 240 RPM
```

Supported Platforms

This change applies to 4220-16P-4XE, 4220-12T-4X, 4220-12P-4X, 4220-4MW-8P-4X, 5320-16P-4XE, and 5320-16P-4XE-DC.

FEC Statistics Detailed Display

Version 34.1.1 adds detailed Forward Error Correction (FEC) statistics for optical interfaces, to help diagnose marginal links, optics issues, and cabling problems.

Overview: FEC statistics provide visibility into the number of corrected and uncorrectable blocks detected on a port, and show which FEC clause is in use. Rising error counts indicate degrading optical signal quality before complete link failure occurs.

Available Statistics:

- Corrected blocks
- Uncorrectable blocks
- FEC clause

Modified CLI Commands

The FEC statistics are added to the output of the existing show ports forward-error-correction command.

View FEC status and statistics:

```
show ports port_list forward-error-correction [no-refresh]
```

Supported Platforms

All platforms with optical interfaces supporting FEC.

FIPS Compliance: Microsoft Golang for OpenSSL Support

Version 34.1.1 updates core components to use Microsoft's Golang distribution to maintain support for the OpenSSL interface required for FIPS compliance.

Overview: Google has discontinued support for a Golang interface to OpenSSL. To ensure the Extreme-labeled FIPS certificate covers all required components, Version 34.1.1 updates telegraf and inlets to use Microsoft's Golang version, which maintains OpenSSL interface support. This change enables continued use of the privately labeled OpenSSL FIPS module and ensures ongoing FIPS compliance certification.

Impact: FIPS-certified deployments maintain cryptographic compliance without requiring customer configuration changes.

Affected Platforms

All platforms with FIPS compliance requirements.

Flood Rate-Limit Bandwidth Mode

Version 34.1.1 extends flood rate-limiting to support bandwidth-based rate limiting in addition to packet-count limiting, enabling byte-rate control for flooded traffic.

Overview: Flood rate-limiting accepts bandwidth rate limits in addition to packet-per-second (pps) limits. Specify the rate in Kbps, Mbps, or Gbps, with an optional burst size, for broadcast, multicast, and unknown destination MAC traffic.

Modified CLI Commands

Configure flood rate-limit in bandwidth mode:

configure ports [*port_list* | *port_group*] **rate-limit flood** [**broadcast** | **multicast** | **unknown-destmac**] [**no-limit** | *pps* | *flood_bandwidth_rate* [**Kbps** | **Mbps** | **Gbps**] {**flood-burst-size** *burst_size* [**Kb** | **Mb**]}]

Example: The following command limits broadcast flood traffic on port 1 to 200 Mbps:

```
configure ports 1 rate-limit flood broadcast 200000 Kbps flood-burst-size 100 Kb out-
actions log disable-port
```

Monitoring

View the flood rate-limit configuration:

show ports *port_list* **information detail**

The following example shows the flood rate-limit settings in the output for port 1:

```
show ports 1 information detail
        Egress Port Rate:       No-limit
        Broadcast Rate:         200000 Kbps, Burst Size: 100 Kb
        Multicast Rate:         400000 Kbps, Burst Size: 200 Kb
        Unknown Dest Mac Rate:  500000 Kbps, Burst Size: 500 Kb
```

Supported Platforms

All platforms.

Hybrid PQC Support in RADIUS TLS and Syslog TLS

Version 34.1.1 adds support for hybrid post-quantum cryptography key encapsulation mechanism (KEM) X25519MLKEM768 in RADIUS TLS and syslog TLS connections. OpenSSL is upgraded from 3.0.10 to 3.5.8 (LTS version) to support this hybrid algorithm, and the maximum TLS version is now TLS 1.3. Support for RSA PSS signature algorithms (rsa_pss_rsae_sha256, rsa_pss_rsae_sha384, rsa_pss_rsae_sha512) is also added.

This feature is automatically enabled and requires no configuration. The switch automatically negotiates X25519MLKEM768 during TLS handshake with compatible RADIUS TLS and syslog TLS servers. The RADIUS TLS or syslog TLS server must support TLS 1.3 and OpenSSL 3.5.0 or later for the hybrid PQC algorithm to be used. If the server does not support this algorithm, the connection proceeds with standard TLS negotiation. Verify server support by checking

openssl version and available TLS groups using `openssl list -kem-algorithms` and `openssl list -tls-groups` commands.

Verification

Verify RADIUS TLS or syslog TLS server compatibility:

Check server OpenSSL version:

```
openssl version
```

List available KEM algorithms:

```
openssl list -kem-algorithms
```

List supported TLS groups:

```
openssl list -tls-groups
```

To verify X25519MLKEM768 negotiation with the RADIUS TLS or syslog TLS server, collect a packet capture and examine the TLS handshake. Look for supported_groups extension with value 0x11ec (X25519MLKEM768) and key_share extension with 1216-byte key exchange data (X25519: 32 bytes + MLKEM768: 1184 bytes).

Supported Platforms

All platforms.

MAC Locking Last Violating MAC Address Clear

Version 34.1.1 adds the ability to clear the last violating MAC address entries from MAC locking violation records without requiring a full system reboot.

Overview: The MAC locking feature restricts port access to specific MAC addresses. When a security violation occurs, the switch increments a violation counter and records the "Last Violating MAC Address" for the affected port. Previously, these records persisted until system reboot, making it difficult for administrators to determine if a violation was historical or recent. Version 34.1.1 allows direct clearing of these records on specific ports.

New CLI Commands

Clear last violating MAC address from specified ports:

`clear mac-locking last-violating-mac ports` *`port_list`*

Supported Platforms

All platforms with MAC locking support.

Network Zones in ACL Policy Rules

Version 34.1.1 extends ACL policy rules to support both source and destination network zone matching in a single rule, and introduces live refresh of zone membership changes without policy rebinding.

Issue: Previous versions supported only source-zone or destination-zone matching in ACL rules. Policies requiring zone-based restrictions on both source and destination required multiple overlapping rules, reducing policy clarity and increasing management complexity. Additionally, zone membership changes required policy rebinding to take effect.

Resolution: ACL policy rules now accept both `source-zone` and `destination-zone` match conditions in a single rule, expanding to the cross product of zone members. The same zone can be used as both source and destination. Zone membership changes take effect immediately using the refresh command without unbinding policies.

Key Capabilities:

- Single ACL rules match packets based on both source and destination network zones (cross-product expansion: M source members × N destination members)
- Same network zone can be used as both source and destination to match all member-to-member pairs
- Zone membership changes (add/delete members) take effect immediately without policy rebinding
- Zone-based policies replace address-based ACLs for improved maintainability and security
- Supports all standard ACL actions with zone-based matching

Modified CLI Commands

Configure an ACL policy with dual-zone (source and destination) matching:

entry *entry_name* { **if** { **source-zone** *source_zone*; **destination-zone** *destination_zone*; } **then** { **deny** | **permit**; } }

Example: The following example counts traffic from zone1 to zone2. Each zone has two members, so the rule expands to four rules (2 source members × 2 destination members).

Create the zones and add their members:

```
create access-list network-zone zone1
create access-list network-zone zone2
configure access-list network-zone zone1 add ipaddress 10.10.10.1 255.255.255.0
configure access-list network-zone zone1 add ipaddress 20.10.10.1 255.255.255.0
configure access-list network-zone zone2 add ipaddress 30.10.10.5 255.255.255.0
configure access-list network-zone zone2 add ipaddress 40.10.10.5 255.255.255.0
```

Create the policy and its rule:

```
edit policy zonepolicy
entry entry1 {
  if {
    source-zone zone1;
    destination-zone zone2;
```

```
  } then {
    count entry1Count;
  }
}
```

Apply the policy to port 10:

```
configure access-list zonepolicy port 10
```

Display the rules that the zone rule expanded to:

```
show access-list port 10 detail
#Policy Entries
#   RuleNo
#   1
entry Cl0:0_30.10.10.5_E1 { if match all {
    destination-address 30.10.10.5/255.255.255.0 ;
    source-address 10.10.10.1/255.255.255.0 ;
} then {
    count entry1Count ;
} }
#   RuleNo
#   2
entry Cl0:0_40.10.10.5_E2 { if match all {
    destination-address 40.10.10.5/255.255.255.0 ;
    source-address 10.10.10.1/255.255.255.0 ;
} then {
    count entry1Count ;
} }
#   RuleNo
#   3
entry Cl0:0_30.10.10.5_E3 { if match all {
    destination-address 30.10.10.5/255.255.255.0 ;
    source-address 20.10.10.1/255.255.255.0 ;
} then {
    count entry1Count ;
} }
#   RuleNo
#   4
entry Cl0:0_40.10.10.5_E4 { if match all {
    destination-address 40.10.10.5/255.255.255.0 ;
    source-address 20.10.10.1/255.255.255.0 ;
} then {
    count entry1Count ;
} }
```

Example: The following rule blocks all traffic among the members of a single zone, by using the same zone as both source and destination:

```
entry block_internal { if { source-zone internal; destination-zone internal; } then
{ deny; } }
```

Live refresh zone membership for all bound policies:

`refresh access-list network-zone`

Zone members can be added or deleted using existing network-zone commands, then applied immediately with refresh without unbinding policies.

Supported Platforms

All platforms with ACL and policy support. Hardware ACL capacity by platform limits zone sizes usable in dual-zone rules.

OpenLDAP Upgrade to 2.6.12

Version 34.1.1 upgrades OpenLDAP from 2.4.21 to 2.6.12.

Existing LDAP server configurations continue to function without modification. Verify LDAP server compatibility with OpenLDAP 2.6.12 in test environments before upgrading production switches.

Supported Platforms

All platforms.

NetLogin Client IP Tracking for RADIUS Framed-IP-Address

Version 34.1.1 improves how NetLogin learns and tracks the IP address of an authenticated client. It fills in the RADIUS Framed-IP-Address attribute and tells the RADIUS server whenever the client IP address changes. This works whether or not the client VLAN has an IP address configured.

Overview: NetLogin learns client IP addresses by listening for events from two sources: the node alias module and ARP.

Node alias events: The node alias module notifies NetLogin when:

- A new client with an IP address is added to the node alias table.
- The IP address of an existing client changes.

The node alias table is updated from the traffic the switch receives from the client.

ARP events: When the client VLAN has an IP address configured, the switch adds the client IP address to the ARP table when it receives traffic from the client. NetLogin also receives these ARP updates.

The events that NetLogin relies on depend on how the client VLAN is configured:

- **Client VLAN without an IP address configured (Layer 2 only):** The ARP table is not populated for the client. NetLogin learns the client IP address only from node alias events.
- **Client VLAN with an IP address configured:** NetLogin learns the client IP address from both node alias events and ARP events. Either one can report the IP address first, and NetLogin uses whichever arrives first.

RADIUS Accounting Behavior

When NetLogin receives a client IP address from a node alias or ARP event, it compares that address with the one it already has for the client.

- When the IP address is learned for the first time, NetLogin sends an Accounting Interim-Update to the RADIUS server with the Framed-IP-Address attribute.
- After that, NetLogin sends the Framed-IP-Address in an Accounting Interim-Update only when the client IP address changes.
- If the reported IP address is the same as the stored address, no Interim-Update is sent. This avoids redundant accounting messages to the RADIUS server.

Configuration

Enable node alias on the ports where NetLogin clients are connected:

enable nodealias ports *port_list*

Node alias learns entries from client traffic only on ports where it is enabled. After it learns an entry, it notifies the NetLogin module of the client IP address.

Supported Platforms

All platforms with NetLogin and node alias support.

Power Supply Failure Alert Severity Change

Version 34.1.1 changes the syslog alert severity level for power supply failures in multi-PSU systems from Info to Warning.

Overview: In multi-PSU systems, when a power supply unit fails, the system now generates a Warning-level syslog message instead of an Info-level message. This change ensures that PSU failures receive appropriate priority in monitoring and alerting systems.

Impact: Monitoring systems configured to trigger alerts on Warning-level messages will now detect PSU failures. Existing monitoring rules that filter on Info-level messages may need adjustment.

Affected Platforms

All platforms with multiple power supply units.

PQC Algorithms in OpenSSH

Version 34.1.1 integrates post-quantum cryptography (PQC) algorithms into OpenSSH, providing hybrid key exchange methods that resist attacks from quantum computers.

Overview: OpenSSH now supports hybrid post-quantum key exchange algorithms including ML-KEM-768 combined with X25519. These hybrid algorithms maintain compatibility with classical cryptography while adding quantum resistance. The algorithm is used only when the

remote end runs OpenSSH 9.5 or later, or a switch running 34.1.1 or later. Otherwise the session falls back to classical key exchange.

New CLI Commands



Note
By default PQC algorithm is enabled on version 34.1.1.

Enable post-quantum key exchange algorithm:

```
configure ssh2 enable post-quantum kex-algorithm mlkem768x25519-sha256
```

Disable post-quantum key exchange algorithm:

```
configure ssh2 disable post-quantum kex-algorithm mlkem768x25519-sha256
```

Monitoring

View SSH key exchange configuration:

```
show ssh2
```

Supported Platforms

All platforms.

SNMP Interface Pause Frame Counters

Version 34.1.1 adds SNMP MIB support for IEEE 802.3x pause frame counters, enabling network management systems to monitor flow control statistics on physical interfaces.

Overview: The switch implements the `dot3PauseTable` SNMP table from RFC 3635, which provides counters and statistics for IEEE 802.3x flow control (pause) frames transmitted and received on each physical interface. This enables operators to monitor flow control behavior and detect congestion patterns.

Available Counters:

- `dot3PauseAdminMode`: Configured pause frame behavior on the interface
- `dot3PauseOperMode`: Actual operational pause frame behavior
- `dot3InPauseFrames`: Count of pause frames received
- `dot3OutPauseFrames`: Count of pause frames transmitted
- `dot3HCInPauseFrames`: 64-bit count of pause frames received
- `dot3HCOutPauseFrames`: 64-bit count of pause frames transmitted

New CLI Commands

No configuration required. Pause frame counters are automatically available through SNMP queries.

Query pause frame statistics by using SNMP:

```
snmpget -v3 -u username switch_ip \ dot3InPauseFrames.ifindex \
dot3OutPauseFrames.ifindex
```

View pause frame statistics:

```
show flow-control statistics ports port_list
```

Supported Platforms

All platforms with flow control support.

SNMPv3 User Authentication with SHA-2

Version 34.1.1 adds support for SHA-256, SHA-384, and SHA-512 HMAC authentication algorithms in SNMPv3 user authentication, replacing SHA-1 for enhanced security.

Issue: SNMPv3 users were limited to SHA-1 HMAC authentication, which does not meet current security requirements for cryptographic hash functions.

Resolution: The switch now accepts SHA-2 family algorithms (SHA-256, SHA-384, SHA-512) for SNMPv3 user authentication. Users can specify the desired algorithm during user configuration.

Supported Algorithms:

- SHA-1 (for backward compatibility)
- SHA-256 (recommended minimum)
- SHA-384
- SHA-512

Modified CLI Commands

Create SNMPv3 user with SHA-256 authentication:

```
configure snmpv3 user username authentication sha-256 password password
```

Create SNMPv3 user with SHA-512 authentication and privacy:

```
configure snmpv3 user username authentication sha-512 password password
privacy aes-256 privacy-password password
```

View SNMPv3 user authentication algorithm:

```
show snmpv3 user username
```

Migration from SHA-1

Existing SNMPv3 users configured with SHA-1 continue to function. Migrate to SHA-256 or higher by reconfiguring users with the new algorithm and updating NMS systems to use the new algorithm for authentication.

Supported Platforms

All platforms.

Transparent Clock Support on 5120 Series

Version 34.1.1 adds Transparent Clock (TC) support on 5120 Series switches, enabling precise time distribution in IEEE 1588 Precision Time Protocol (PTP) deployments.

Overview: Transparent Clock mode allows the switch to participate in PTP timing distribution without maintaining a local system clock. The switch measures the transit time for PTP packets crossing its switching fabric and updates the `correctionField` in each packet to compensate for the delay introduced by the switch.

Clock Modes:

- **End-to-End TC:** Measures the time PTP packets spend between ingress and egress ports

CLI Commands

Create the end-to-end transparent clock:

```
create network-clock ptp end-to-end-transparent
```

Add ports to the end-to-end transparent clock:

```
configure network-clock ptp end-to-end-transparent add ports port_list
one-step
```

Enable the end-to-end transparent clock on ports:

```
enable network-clock ptp end-to-end-transparent ports port_list
```

Monitoring

Display the end-to-end transparent clock:

```
show network-clock ptp end-to-end-transparent
```

The following is sample output:

```
show network-clock ptp end-to-end-transparent

================================================================================
                        Transparent Clock Information
================================================================================
Enabled               :  TRUE
Clock Mode            :  End-to-end
```

```
Clock Instance          :  1
Clock Domain            :  0
================================================================================
```

Supported Platforms

5120 Series only.

Update Time-Zone Information for UTC Changes in Alberta and British Columbia

Version 34.1.1 updates Switch Engine to support time-zone changes in Canada. As of November 2026, the time zone for Alberta permanently stays at UTC-6. The same applies to the Northwest Territories. British Columbia permanently stays at UTC-7.

Overview: These provinces and territories keep a single UTC offset all year, so automatic Daylight Saving Time (DST) changes no longer apply to them.

Configuration

Specify the GMT offset in minutes and disable automatic DST.

Configure the time zone for Alberta or the Northwest Territories (UTC-6):

```
configure timezone -360 noautodst
```

Configure the time zone for British Columbia (UTC-7):

```
configure timezone -420 noautodst
```

Supported Platforms

All platforms.

Updates to Avoid NTP Epocalypse - Y2036 Overflow

Version 34.1.1 updates the operating system (OS) so that Network Time Protocol (NTP) continues to keep correct time after the NTP era rollover in 2036.

Overview: NTP timestamps count seconds since January 1, 1900, in a 32-bit field. The field overflows on February 7, 2036, which can cause systems that rely on it to misread the time. Similar to the Y2038 issue, the OS was investigated for this vulnerability, and the fixes are included in this release.

Supported Platforms

All platforms.

Updates to Avoid Epocalypse - Y2038 Overflow

Version 34.1.1 updates the software to use 64-bit integer time values, avoiding the Year 2038 overflow (also known as Y2038 or the Epochalypse).

Overview: Many systems track time as the number of seconds since January 1, 1970, at 00:00:00 UTC. On older systems, this number is stored in a signed 32-bit integer, which can hold a maximum of 2,147,483,647 seconds. At 03:14:08 UTC on January 19, 2038, the next second exceeds this limit, and the value wraps around to a negative number, which the system interprets as a time in December 1901. All components that use the internal time variable now use a 64-bit value to avoid the overflow.

Supported Platforms

All platforms.

Changing the Network Operating System

Universal Hardware switches can run two different operating systems: Switch Engine (default) or Fabric Engine.

Making Your Initial Network Operating System Selection

You can make your initial selection of the operating system using:

- **ExtremeCloud IQ** —You can select your network operating system when purchasing your switch, which associates the switch serial number with your desired network operating system, which then causes the desired network operating system to be loaded during ExtremeCloud onboarding. For more information about using ExtremeCloud IQ, go to https://www.extremenetworks.com/support/documentation/extremecloud-iq/.
- **Extreme Management Center**— see documentation for version 22.3 or later
- **Manually during boot-up**:
  - **Bootloader**—When you see the message `Starting Default Bootloader ...Press and hold the <spacebar> to enter the bootrom`, press and hold the **space bar** until the boot menu is displayed (you have 30 seconds):

```
*** 5320-48T-8XE Boot Menu ( 3.4.2.8 ) ***

     EXOS: Default
     EXOS: Primary 32.1.1.6
     EXOS: Secondary 32.1.1.6
     EXOS: Primary 32.1.1.6 with default configuration
     EXOS: Secondary 32.1.1.6 with default configuration
     EXOS: Rescue
     Change the switch OS to VOSS
     Run Manufacturing Diagnostics
     Update bootloader
     Reboot system
```

    Use the **up** and **down** arrow keys to select `Change the switch OS to VOSS`, and then press **Enter**.

    Note
    The 5720, 7520, and 7720 Series use the **GRUB** menu. There is no need to press and hold the **space bar**. Use the **up** and **down** arrow keys to navigate the menu.

  - **Safe defaults mode start-up menu**—When the question `Would you like to change the switch OS to VOSS? [y/N/q]` is displayed:
    - For Switch Engine, type `N`.
    - For Fabric Engine, type `y`.

Continue to log onto the switch. For more information about logging onto the switch, see the user guide.

Changing Your Network Operating System

You can change your network operating system selection at any time.



Caution
Changing your network operating systems deletes all configuration files, debug information, logs, events, and statistics information of the previous network operating system.



Note
If you anticipate ever changing the operating system to Fabric Engine, and you want to statically assign IP addresses on the DHCP server, then it is recommended to assign them based on the DHCP client ID. For more information about this issue, see the *Using a BOOTP or DHCP Server* topic in the user guide.

- **ExtremeCloud IQ**—See https://www.extremenetworks.com/support/documentation/extremecloud-iq/
- **Extreme Management Center**—See *Extreme Management Center User Guide*
- **CLI Command**—run the `download [url` *`url`* `{vr` *`vrname`*`} | image [active | inactive] [[`*`hostname | ipaddress`*`] `*`filename`*` {{vr} `*`vrname`*`} {block-size `*`block_size`*`}] {`*`partition`*`} {install {reboot}}` command specifying a VOSS image.

  Note
  Do *not* use the active, inactive, and partition options. They are not applicable for Fabric Engine.

ExtremeCloud IQ Agent Support

Switch Engine supports ExtremeCloud IQ. For network administrators looking for unified management of access points, switches, & routers, ExtremeCloud IQ is a cloud-driven network management application that:

- simplifies network operations through an easy to use and intuitive interface, including minimal touch onboarding of devices
- provides ultimate flexibility in deployment choice, cloud platform choice, OS choice
- offers unlimited data duration for more informed networking decisions

Important
Check the ExtremeCloud IQ release notes to ensure support for your version has been added before upgrading.

This version supports device discovery, basic monitoring, visibility into homogenous stacking, and the ability to configure an optional user-defined virtual router (VR) and address of the server for ExtremeCloud IQ agent to connect to. These values are used instead of any auto-detected values.

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

Table 6: Supported Platforms

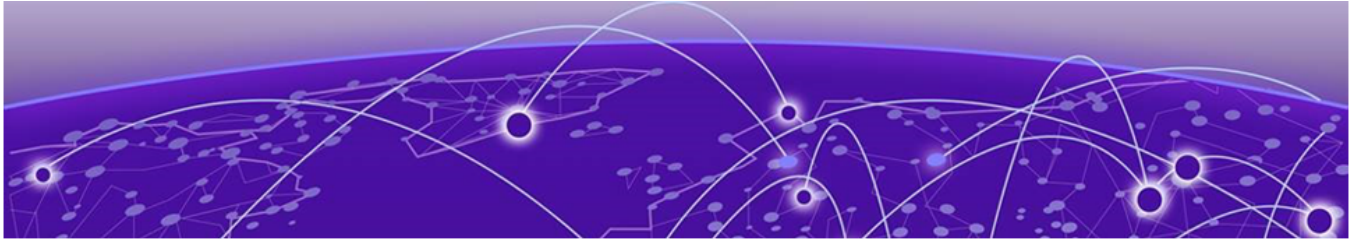
Switch Series	Switch Models
4120	4120-24MW-4Y 4120-48MW-4Y
4220	4220-8X 4220-12P-4X 4220-12T-4X 4220-24P-4X 4220-24T-4X 4220-48P-4X 4220-48T-4X 4220-4MW-8P-4X 4220-4MW-20P-4X 4220-8MW-40P-4X
5120	5120-24X-4Y 5120-24XT-4Y 5120-44X-4Y-2C

Table 6: Supported Platforms (continued)

Switch Series	Switch Models
5320	5320-48T-8XE 5320-48P-8XE 5320-24T-8XE 5320-24P-8XE 5320-16P-4XE 5320-16P-4XE-DC 5320-24T-4X-XT 5320-24T-24S-4XE-XT
5420	5420F-8W-16P-4XE 5420F-24P-4XE 5420F-24S-4XE 5420F-24T-4XE 5420F-16MW-32P-4XE 5420F-16W-32P-4XE 5420F-48P-4XE 5420F-48P-4XL 5420F-48T-4XE 5420M-24T-4YE 5420M-24W-4YE 5420M-16MW-32P-4YE 5420M-24W-24S-4YE 5420M-48T-4YE 5420M-48W-4YE
5520	5520-24T 5520-24W 5520-48T 5520-48W 5520-12MW-36W 5520-24X 5520-48SE 5520-24T-ACDC-BASE 5520-48T-ACDC-BASE 5520-24X-ACDC-BASE 5520-48SE-ACDC-BASE
5720	5720-24MW 5720-24MXW 5720-48MW 5720-48MXW

Table 6: Supported Platforms (continued)

Switch Series	Switch Models
7520	7520-48Y-8C 7520-48XT-6C 7520-48YE-8CE
7720	7720-32C



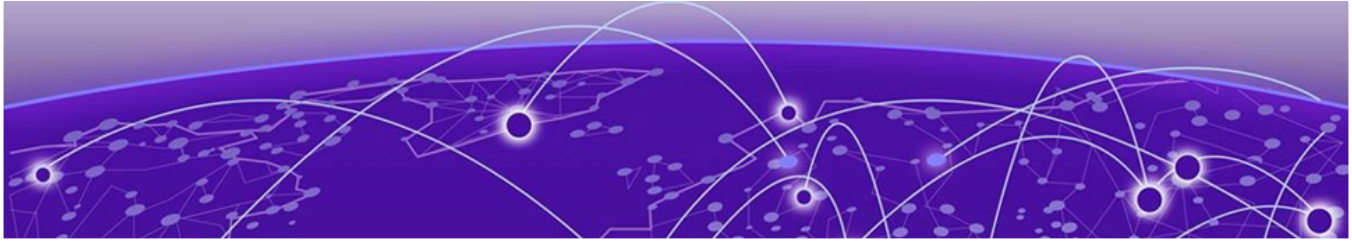
Extreme Hardware/Software Compatibility and Recommendation Matrices

ExtremeXOS and Switch Engine Software Support provides information about the minimum version of software required to support switches.

The Extreme Optics Compatibility website displays supported hardware platforms, technical specifications, and usage considerations for pluggable optical devices (transceivers and cables) used in all Extreme Networks operating environments. To access the site, open https://optics.extremenetworks.com/EXOS/ in a web browser.

To find the recommended releases for Universal Hardware platforms, see *ExtremeXOS and Switch Engine Release Recommendations*.

The latest versions of this and other guides are at: www.extremenetworks.com/documentation/.

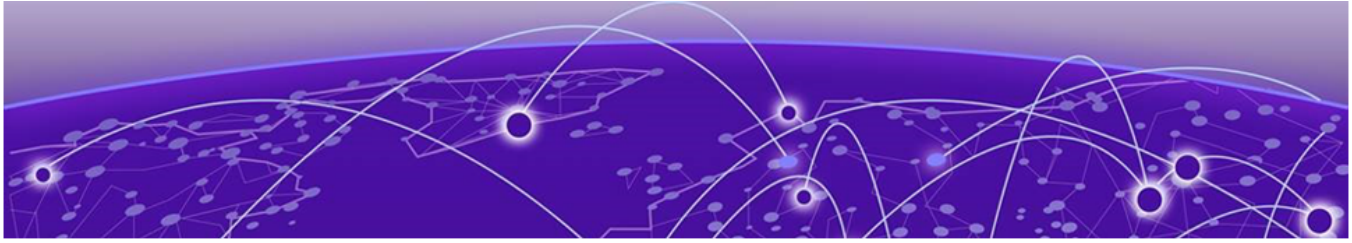


Compatibility with Extreme Management Center

This version of Switch Engine is compatible with the version of Extreme Management Center as shown in this table: http://emc.extremenetworks.com/content/common/releasenotes/extended_firmware_support.htm.

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

This version was tested with ExtremeCloud IQ Site Engine version 25.11.10.48.

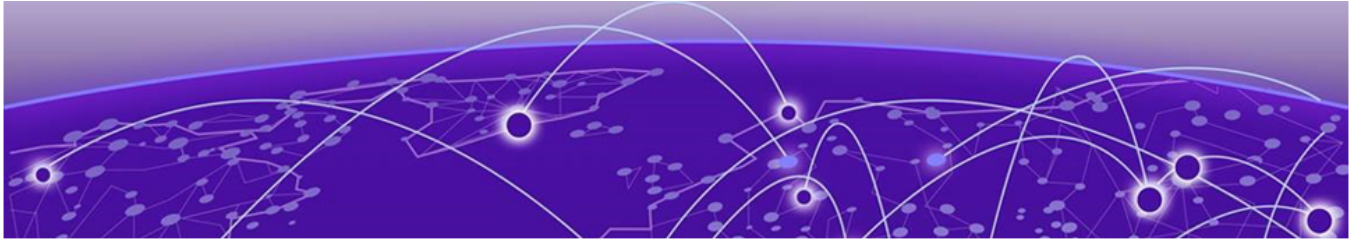


Supported MIBs

The Extreme Networks management information bases (MIBs) are located on the Extreme Portal in the Downloads section. Log in to the Extreme Portal to view and download.

When you provide your serial number or agreement number, the MIBs are available under each release.

For detailed information on which MIBs and SNMP traps are supported, see the *Extreme Networks Proprietary MIBs* and *MIB Support Details* sections in the user guide.



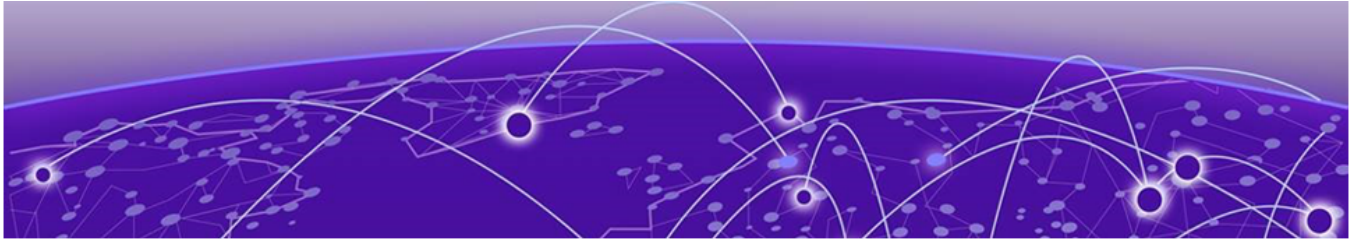
Tested Third-Party Products

The following third-party products have been tested.

Tested RADIUS Servers

The following RADIUS servers are fully tested:

- Microsoft—Internet Authentication Server
- Meetinghouse
- FreeRADIUS



Extreme Switch Security Assessment

DoS Attack Assessment

Tools used to assess DoS attack vulnerability:

- Network Mapper (NMAP)

ICMP Attack Assessment

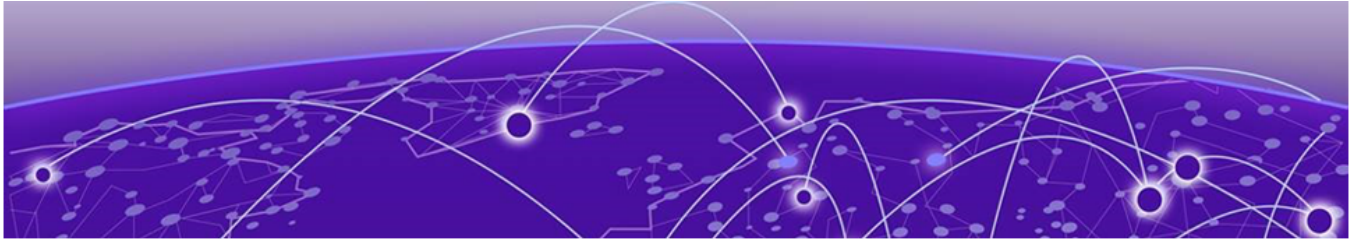
Tools used to assess ICMP attack vulnerability:

- SSPing
- Twinge
- Nuke
- WinFreeze

Port Scan Assessment

Tools used to assess port scan assessment:

- Nessus



Limits

Limits Overview on page 53
Base License Limits on page 56
Premier License Limits on page 93
Notes for Limits Tables on page 103

This chapter summarizes the supported limits in this version.

Limits Overview

The limits data is grouped by license level that contains the associated features:

- Base License Limits on page 56
- Premier License Limits on page 93

The Universal family of switches includes two license levels: Base and Premier.

The following figure illustrates that each license level builds on the features of the license level below it. For example, the Premier license includes all of the features in the Base license, plus the features in the Premier license level.

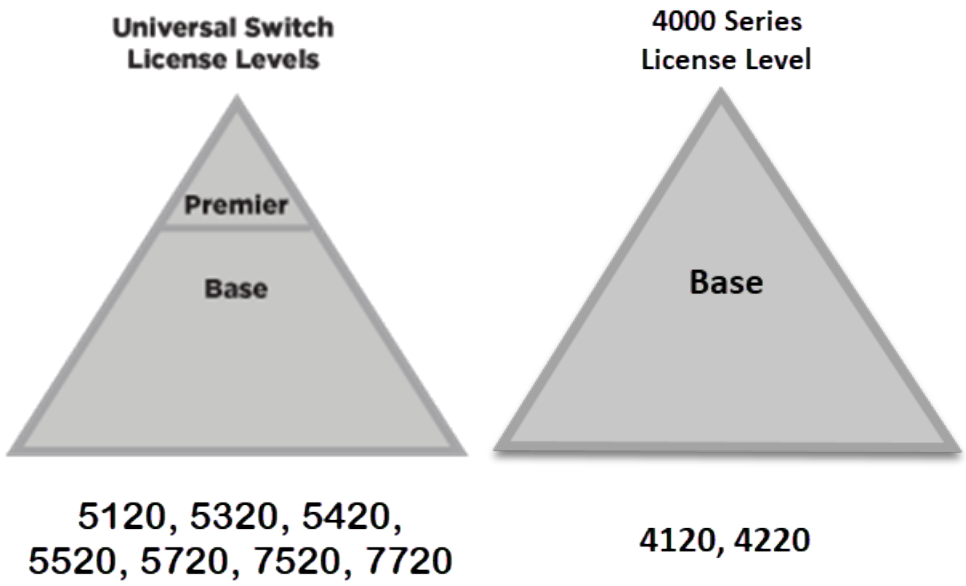


Figure 1: License Levels for Universal Switches

Extreme Platform ONE includes three license levels: Standard, Advanced, and Premium. A Standard license is required to manage devices from ExtremeCloud IQ.

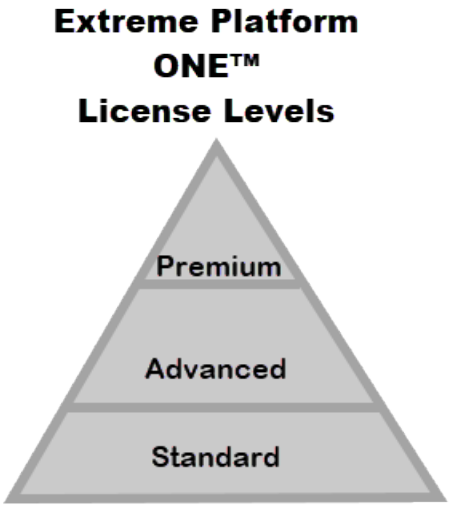


Figure 2: Extreme Platform ONE License Levels

Each license level is purchased based on four tiers, depending on device type:

- A - 4000 series, 5120, 5320
- B - 5420

- C - 5520
- D - 5720, 7520, 7720

Universal devices with a verified Extreme Platform ONE license will perform the following actions:

- 5000 and 7000 series - activate Premier Universal license features

Extreme Platform ONE also provides operating system product service, management, and insights.

For more information about licenses, see Switch Engine v33.6.1 Licensing Guide .

The following tables summarize tested metrics for a variety of features, as measured in a per-system basis unless otherwise noted. These limits may change, but represent the current status. The contents of this table supersede any values mentioned in the Switch Engine books.

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.

The route limits shown in the following tables for IPv4 and IPv6 routing protocols are software limits only. The actual hardware limits may be higher or lower than the software limits, based on platform. The hardware limits for specific platforms are specified as "IPv4/IPv6 routes (LPM entries in hardware)" in the following tables.

In the architecture, Layer-2, Layer-3, and multicast packet forwarding and filtering operations take place on the controlling bridge. The controlling bridge switch and attached BPEs (V400 Virtual Port Extenders) constitute a single, extended switch system. Therefore, the system assumes the scale and limits from the specific controlling bridge model in use. For applicable limits, see the following tables for the controlling bridge you are using.

Base License Limits

The following table shows supported limits for features in the Base License.

Table 7: Supported Limits for the Base License

Metric	Product	Limit
AAA (local)—maximum number of admin and local user accounts.	All platforms	16
Access lists (meters)—maximum number of meters.	4120, 5120	512 ingress 128 egress
	4220	2,048 ingress 256 egress
	5320, 5420	6,144 ingress 512 egress
	5320-16P-2MXT-2X	1,024 ingress 256 egress
	7520, 7720	6,144 ingress 1,024 egress
	5520	2,048 ingress 512 egress
	5720-MW	6,144 ingress 3,072 egress
	5720-MXW	6,144 ingress 6,144 egress
Access lists (policies)—suggested maximum number of lines in a single policy file.	All platforms	300,000

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Access lists (policies)—maximum number of rules in a single policy file. [a]	4220, 5320-48T/P, 7520, 7720	8,192 ingress 1,024 egress
	5320-24T/P, 5320-16P	8,192 ingress 512 egress
	5320-16P-2MXT-2X	1,000 (rules double-wide (160-bit)) ingress 2,000 (rules single-wide (80-bit, default)) ingress 512 egress
	4120, 5120	1,024 ingress 256 egress
	5420M	18,000 (rules double-wide (160-bit)) ingress 36,000 (rules single-wide (80-bit, default)) ingress 1,024 egress
	5420F	8,000 (rules double-wide (160-bit)) ingress 16,000 (rules single-wide (80-bit, default)) ingress 1,024 egress
	5520	9,216 ingress 1,024 egress
	5720-MW	18,432 (80-bit) ingress 6,144 egress
	5720-MXW	36,864 (80-bit), 18,432 (160-bit) ingress 12,288 egress

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Access lists (policies)— maximum number of rules in a single policy file in first stage (VFP).	5520, 5720	2,048 ingress only
	5320-48T/P, 5420, 7520, 7720	1,024 ingress only
	4220, 5320-16P, 5320-24T-4X-XT	512 ingress only
	4120, 5120	256 ingress
Access lists (slices)—number of ACL slices.	5720, 7520, 7720	12 ingress 4 egress
	5320-48T/P, 5420, 5520	18 ingress 4 egress
	4120, 4220, 5120, 5320-24T/P, 5320-16P	8 ingress 4 egress
Access lists (slices)—number of ACL slices in first stage (VFP).	All platforms	4 ingress only
ACL Per Port Meters—number of meters supported per port.	All platforms	16
ACL port ranges.	All platforms	32
Meters Packets-Per-Second Capable.	All platforms	N/A
AVB (audio video bridging)—maximum number of active streams.	5120, 5320, 5420	1,024
	5520, 5720, 7520	4,096
BFD sessions (Software Mode)—maximum number of BFD sessions.	5320, 5420, 5520, 5720, 7520, 7720 (default timers—1 sec).	512
	5120 (default timers—1 sec).	90
BFD IPv4 sessions (Hardware Assisted)—maximum number of IPv4 BFD sessions.	7520, 7720	900 425 256 (with 3 ms transmit interval)
BFD IPv6 sessions (Hardware Assisted)—maximum number of IPv6 BFD sessions.	7520, 7720	425 (PTP not enabled)

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
BGP (multicast address-family routes)—maximum number of multicast address-family routes.	5520, 5720-MXW	13,000
	5720-MW	20,000
	7520, 7720	25,000
	5320-16P-4XE, 5320 24-port except XT	8,000
	5320 48-port, 5420	12,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	992
	5120	64
BGP (non-unique routes)—maximum number of nonunique BGP routes.	7520. 7720	75,000
	5720-MW	60,000
	5320 48-port, 5420	36,000
	5320-16P-4XE, 5320 24-port except XT	24,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	2,700
	5120	192
BGP (peers)—maximum number of BGP peers.	All platforms except 4120 and 4220	2
BGP (unicast address-family routes)—maximum number of unicast address-family routes.	5520, 5720-MW (at default)	13,000
	5720-MXW (at default)	20,000
	7520, 7720 (at deafult)	25,000
	5320 48-port, 5420	12,000
	5320-16P-4XE, 5320 24-port except XT	8,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	992
	5120	64
	5720-MW (with ALPM enabled)	163,000
	5720-MXW (with ALPM enabled)	288,000
	5520 (with ALPM enabled)	80,000
BGP auto-peering—maximum number of auto-peering nodes and VTEPs.	All platforms except 4120 and 4220	64
BGP auto-peering attached IPv4 hosts— maximum number of attached IPv4 hosts.	All platforms except 4120 and 4220	64,000
BGP auto-peering attached IPv6 hosts— maximum number of attached IPv6 hosts.	All platforms except 4120 and 4220	8,000

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
BGP auto-peering ECMP—maximum number of equal cost multipath for auto-peering. **Note:** * Subject to the limitation imposed by the number of physical ports on a switch.	5720, 7520, 7720 5320, 5420, 5520	16* 4*
BGP auto-peering maximum IPv4 prefixes with ECMP—Maximum number of IPv4 Network prefixes with ECMP.	5120, 5320, 5420, 5520, 5720 7520, 7720	16,000 64,000
BGP auto-peering maximum IPv6 prefixes with ECMP—Maximum number of IPv6 Network prefixes with ECMP.	5120, 5320, 5420, 5520, 5720 7520, 7720	254 64,000
BGP auto-peering MLAG peers—maximum MLAG peers per AutoBGP node.	All platforms except 4120 and 4220	1
BGP auto-peering VRFs—maximum number of VRFs.	All platforms except 4120 and 4220	64
BGP auto-peering EVPN instances—maximum EVPN instances.	All platforms except 4120, 4220, and 5120	1,024
BGPv6 (unicast address family routes)—maximum number of unicast address family routes.	5320 48-port, 5420, 5520, 5720-MW (at default)	6,000
	5720-MW (with ALPM enabled)	107,000
	5720-MXW, 7520, 7720 (at default)	10,000
	5120	64
	5720-MXW (with ALPM enabled)	213,000
	5520 (with ALPM enabled)	40,000
	5320-16P-4XE, 5320 24-port except XT	4,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	496
BGPv6 (non-unique routes)—maximum number of nonunique BGP routes.	5320 48-port, 5420, 5520, 5720-MW	18,000
	5720-MXW, 7520, 7720	30,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	14,000
	5320-16P-4XE, 5320 24-port except XT	12,000
	5120	64
BOOTP/DHCP relay—maximum number of BOOTP or DHCP servers per virtual router.	All platforms	8

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
BOOTP/DHCP relay—maximum number of BOOTP or DHCP servers per VLAN.	All platforms	8
BOOTP/DHCP relay—maximum number of DHCPv4/v6 relay agents	All platforms	4,000
Connectivity fault management (CFM)—maximum number or CFM domains.	All platforms	8
CFM—maximum number of CFM associations.	All platforms	256
CFM—maximum number of CFM up end points.	All platforms	32
CFM—maximum number of CFM down end points.	All platforms	32
CFM—maximum number of CFM remote end points per up/down end point.	All platforms	2,000
CFM—maximum number of dot1ag ports.	All platforms	128
CFM—maximum number of CFM segments.	All platforms	1,000
CFM—maximum number of MIPs.	All platforms	256
CLEAR-Flow—total number of rules supported. The ACL rules plus CLEAR-Flow rules must be less than the total number of supported ACLs.	4120, 4220, 5120, 5320, 5420, 5720, 7520, 7720 ExtremeSwitching 5520	8,192 9,215
Data Center Bridging eXchange (DCBX) protocol Type Length Value (TLVs)—maximum number of DCBX application TLVs.	All platforms	8
DHCPv6 Prefix Delegation Snooping—Maximum number of DHCPv6 prefix delegation snooped entries.	All platforms	256 (with underlying protocol RIPng) 128 (with underlying protocol OSPFv3) 1,024 (with static routes)

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
DHCP snooping entries—maximum number of DHCP snooping entries.	All platforms	2,048
Dynamic ACLs—maximum number of ACLs processed per second. **Note:** Limits are load-dependent.	All platforms with 50 DACLs with 500 DACLs	 10 5
EAPS domains—maximum number of EAPS domains. **Note:** An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains.	5720 4120, 4220, 5120, 5320-24T/P, 5320-16P 5320-48T/P, 5420, 5520	128 32 64
EAPSv1 protected VLANs—maximum number of protected VLANs.	4120, 4220, 5120, 5320-24T/P, 5320-16P 5320-48T/P, 5420, 5520, 5720, 7520, 7720	1,000 2,000
EAPSv2 protected VLANs—maximum number of protected VLANs.	4120, 4220, 5120, 5320, 5420, 5520 5720, 7520, 7720	1,000 2,000
ELSM (vlan-ports)—maximum number of VLAN ports.	4120, 4220, 5120, 5320-24T/P, 5320-16P 5320-48T/P, 5420, 5520, 5720, 7520, 7720	4,000 5,000
ERPS domains—maximum number of ERPS domains with or without CFM configured.	All platforms	32
ERPSv1 protected VLANs—maximum number of protected VLANs.	4120, 4220, 5120, 5320-24T/P, 5320-16P 5320-48T/P, 5420, 5520, 5720, 7520, 7720	1,000 2,000
ERPSv2 protected VLANs—maximum number of protected VLANs.	4120, 4220, 5120, 5320-24T/P, 5320-16P 5320-48T/P, 5420, 5520, 5720, 7520, 7720	500 2,000
ESRP groups—maximum number of ESRP groups	All platforms	32
ESRP domains—maximum number of ESRP domains.	4220, 5320, 5420, 5520, 5720, 7520, 7720. 4120, 5120	64 32
ESRP L2 VLANs—maximum number of ESRP VLANs without an IP address configured.	4220, 5320, 5420, 5520, 5720, 7520, 7720 4120, 5120	1,000 120
ESRP L3 VLANs—maximum number of ESRP VLANs with an IP address configured.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4220, 5320-24T/P, 5320-16P 4120, 5120	511 509 120

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
ESRP (maximum ping tracks)—maximum number of ping tracks per VLAN.	All platforms	8
ESRP (IP route tracks)—maximum IP route tracks per VLAN.	All platforms	8
ESRP (VLAN tracks)—maximum number of VLAN tracks per VLAN.	All platforms	1
Extended Edge Switching maximum BPEs—maximum number of attached bridge port extenders (BPEs).	5520, 7520-48Y 5420	48 20
Extended Edge Switching maximum cascade ports—maximum number of upstream ports on bridge port extenders (BPEs).	5420, 5520, 7520-48Y	2 on V400-24 and V300 models 4 on V400-48 models
Extended Edge Switching maximum tiers—maximum number of cascade levels (tiers) of bridge port extenders (BPEs).	ExtremeSwitching 5420, 5520, 7520-48Y	4 (except for V300-8P-2T-W, which support 1 tier)
Extended Edge Switching maximum ring BPEs—maximum number of bridge port extenders (BPEs) in a ring topology.	ExtremeSwitching 5420, 5520, 7520-48Y	8
Extended Edge Switching maximum VLANs—maximum number of VLANs - Includes all VLANs	ExtremeSwitching 5520, 7520-48Y ExtremeSwitching 5420	4,094 1,024
Extended Edge Switching VLAN+ port memberships—maximum number of VLAN+ (extended) port memberships.	ExtremeSwitching 5520, 7520-48Y 5420	12,000 in hash mode (default) 131,000 in port-group mode 8,750 in hash mode (default) 131,617 in port-group mode

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Forwarding rate—maximum L3 software forwarding rate.	4220	9,274 pps
	4120	12,624 pps
	5120	9,000 pps
	5320-24P-8XE, 5320-24T-4X-XT	11,000 pps
	5320-48P	19,142 pps
	5420F	21,585 pps
	5520	18,838 pps
	5720-MW	27,000 pps
	5720-MXW	31,000 pps
	7520, 7720	34,813 pps
FDB (unicast blackhole entries)—maximum number of unicast blackhole FDB entries.	4120, 5120	16,384
	4220, 5320	32,000
	ExtremeSwitching 5420M	65,536
	ExtremeSwitching 5420F	32,768 [f]
	ExtremeSwitching 5520	114,688 [f]
	ExtremeSwitching 5720-MW	163,840 [f]
	ExtremeSwitching 5720-MXW, 7520, 7720	294,912 [f]
FDB (multicast blackhole entries)—maximum number of multicast blackhole FDB entries.	5520, 5720-MW	4,096
	4120, 4220, 5120, 5320, 5420	1,024
	5720-MXW, 7520, 7720	16,000
FDB (maximum L2 entries)—maximum number of MAC addresses.	4120, 5120	16,384
	4220, ExtremeSwitching 5320	32,000
	ExtremeSwitching 5420M	65,536
	ExtremeSwitching 5420F	32,768 [g]
	ExtremeSwitching 5520	114,688 [g]
	ExtremeSwitching 5720-MW	163,840 [g]
	5720-MXW, 7520, 7720	294,912 [g]
FDB (maximum L2 entries)—maximum number of multicast FDB entries.	ExtremeSwitching 5520	4,096
	4120, 4220, 5120, 5320, 5420	1,024
	5720, 7520, 7720	16,000
GRE Tunnels—maximum number of GRE tunnels.	All platforms, except 4120, 5120	255
Identity management—maximum number of Blacklist entries.	All platforms except 4120 and 4220.	512

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Identity management—maximum number of Whitelist entries.	All platforms except 4120 and 4220.	512
Identity management—maximum number of roles that can be created.	All platforms except 4120 and 4220.	64
Identity management—maximum role hierarchy depth allowed.	All platforms except 4120 and 4220.	5
Identity management—maximum number of attribute value pairs in a role match criteria.	All platforms except 4120 and 4220.	16
Identity management—maximum number of child roles for a role.	All platforms except 4120 and 4220.	8
Identity management—maximum number of policies/dynamic ACLs that can be configured per role.	All platforms except 4120 and 4220.	8
Identity management—maximum number of LDAP servers that can be configured.	All platforms except 4120 and 4220.	8
Identity management—maximum number of Kerberos servers that can be configured.	All platforms except 4120 and 4220.	20
Identity management—maximum database memory size.	All platforms except 4120 and 4220.	512
Identity management—recommended number of identities per switch. **Note:** Number of identities per switch is for a default identity management database size (512 Kbytes) across all platforms.	All platforms except 4120 and 4220.	100
Identity management—recommended number of ACL entries per identity. **Note:** Number of ACLs per identity, based on system ACL limitation.	All platforms except 4120 and 4220.	20

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Identity management—maximum number of dynamic ACL entries configured as an individual dynamic rule, or as an ACL entry in a policy file.	All platforms except 4120 and 4220.	500
IGMP snooping per VLAN filters—maximum number of VLANs supported in per-VLAN IGMP snooping mode.	ExtremeSwitching 5320 (except 5320-24T-4X-XT), 5420, 5520, 5720, 7520, 7720	1,500
	4220, ExtremeSwitching 5320-24T-4X-XT	500
	4120	48
	5120	100
IGMPv1/v2 SSM-map entries—maximum number of IGMPv1/v2 SSM mapping entries.	5320, 5420, 5520, 5720, 7520, 7720	6
	5120	60
IGMPv1/v2 SSM-map entries—maximum number of sources per group in IGMPv1/v2 SSM mapping entries.	All platforms except 4120 and 4220.	50
IGMPv2 subscriber—maximum number of IGMPv2 subscribers per port. [n]	5320 (except 5320-24T-4X-XT), 5420, 7520, 7720 ,5720,5520	4,000
	4220, 5320-24T-4X-XT	1,000
	4120, 5120	250
IGMPv2 subscriber—maximum number of IGMPv2 subscribers per switch.[n]	ExtremeSwitching 5320 (except 5320-24T-4X-XT), 5420, 5520	20,000
	ExtremeSwitching 5720-MW, 7520, 7720	45,000
	ExtremeSwitching 5720-MXW	54,000
	4220, 5320-24T-4X-XT	1,000
	4120, 5120	256
IGMPv3 maximum source per group—maximum number of source addresses per group.	All platforms	250
IGMPv3 subscriber—maximum number of IGMPv3 subscribers per port. [n]	5320 (except 5320-24T-4X-XT), 5420, 5520, 5720, 7520, 7720	4,000
	4220, 5320-24T-4X-XT	1,000
	4120, 5120	250

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
IGMPv3 subscriber—maximum number of IGMPv3 subscribers per switch. [n]	ExtremeSwitching 5320 (except 5320-24T-4X-XT), 5420, 5520	20,000
	ExtremeSwitching 5720-MW, 7520, 7720	45,000
	ExtremeSwitching 5720-MXW	54,000
	4220, 5320-24T-4X-XT	1,000
	4120, 5120	256
IP ARP entries in software—maximum number of IP ARP entries in software. **Note:** Might be limited by hardware capacity of FDB (maximum L2 entries).	4120, 5120	400
	4220, 5320-16P-2MXT-2X	4,000
	5320 (except 5320-16P-2MXT-2X), 5420F models	12,000
	5420M models	24,000
	5520	74,750 [h]
	5720-MW	100,000
	7520, 7720	184,318 (up to)
	ExtremeSwitching 5720-MXW	221,000
IPv4 ARP entries in hardware with minimum LPM routes—maximum recommended number of IPv4 ARP entries in hardware, with minimum LPM routes present. Assumes number of IP route reserved entries is 100 or less.	4120, 5120	397
	4220	4,000
	5320	12,000
	5320-16P-2MXT-2X	4,000
	ExtremeSwitching 5420M models	24,000
	ExtremeSwitching 5420F models	12,000
	ExtremeSwitching 5520	60,000 [h]
	ExtremeSwitching 5720-MW	80,000 [h]
	7520, 7720	146,000 [h]
	ExtremeSwitching 5720-MXW	172,000 [h]

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
IPv4 ARP entries in hardware with maximum LPM routes—maximum recommended number of IPv4 ARP entries in hardware, with maximum LPM routes present. Assumes number of IP route reserved entries is "maximum."	4120, 5120	384
	4220	3,000
	5320	10,000
	5320-16P-2MXT-2X	3,000
	ExtremeSwitching 5420M models	21,000
	ExtremeSwitching 5420F models	10,000
	ExtremeSwitching 5520	49,000 h
	ExtremeSwitching 5720-MW	70,000 h
	7520, 7720	125,000 h
	ExtremeSwitching 5720-MXW	156,000 h
IP flow information export (IPFIX)—number of simultaneous flows.	ExtremeSwitching 5420	4,000 (IPv4 and IPv6 flows)
	ExtremeSwitching 5520	32,000 (IPv4 flows) 18,000 (IPv6 flows)
	ExtremeSwitching 5720	257,000 (IPv4 flows) 112,000 (IPv6 flows)
IPv4 remote hosts in hardware with zero LPM routes—maximum recommended number of IPv4 remote hosts (hosts reachable through a gateway) in hardware when LPM routing is not used. Assumes number of IP route reserved entries is 0, and number of IPv4 ARP entries present is 100 or less.	4120, 5120	450
	4220	4,000
	5320	20,000
	5320-16P-2MXT-2X	7,000
	ExtremeSwitching 5320-24T/P, 5320-16P	24,000
	ExtremeSwitching 5420M	36,000
	ExtremeSwitching 5420F	24,000 h
	ExtremeSwitching 5520	102,000 h
	ExtremeSwitching 5720-MW	139,000 h
	7520, 7720	241,000 (up to)
	5720-MXW (with ALPM enabled)	245,000 h

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
IPv4 routes—maximum number of IPv4 routes in software (combination of unicast and multicast routes), including static and from all routing protocols.	5520	81,000
	4120, 4220, 5120, 5320, 5420	25,000
	5720-MW	163,000
	5720-MXW	288,000
	7520, 7720	350,000
IPv4 routes (LPM entries in hardware)— number of IPv4 routes in hardware.	4120, 5120	64 [q]
	4220, 5320-16P-2MXT-2X	992
	5320-16T/P, 5320-24T/P	8,000
	5320-48T/P, 5420	12,000
	5520	81,000 [q]
	ExtremeSwitching 5720-MW	163,000 [q]
	7520, 7720	262,000 up to 350,000 [q]
	ExtremeSwitching 5720-MXW	288,000 [q]
IPv6 6in4 tunnel—maximum number of IPv6 6in4 tunnels.	All platforms except 4120, 5120	255
IPv6 6to4 tunnel—maximum number of IPv6 6to4 tunnels.	All platforms except 4120, 5120	1 (per virtual router)
IPv6 addresses on an interface—maximum number of IPv6 addresses on an interface.	All platforms	255
IPv6 addresses on a switch—maximum number of IPv6 addresses on a switch.	All platforms	2,048
IPv6 host entries in hardware—maximum number of IPv6 neighbor entries in hardware.	4120, 5120	200
	4220	2,000
	5320	6,000
	5320-16P-2MXT-2X	3,000
	5420M models	12,000
	ExtremeSwitching 5420F models	6,000
	ExtremeSwitching 5520	18,000 [s]
	ExtremeSwitching 5720-MW	24,000 [s]
	7520, 7720	57,000 [h]
	ExtremeSwitching 5720-MXW	78,000 [s]

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
IPv6 routes in software—maximum number of IPv6 routes in software, including static routes and routes from all routing protocols.	ExtremeSwitching 5520	18,000 [q]
	4120, 4220, 5320, 5420	25,000
	5720-MW	70,000 [q]
	7520, 7720	196,000 [q]
	ExtremeSwitching 5720-MXW	213,000 [q]
IPv6 routes (LPM entries in hardware)—maximum number of IPv6 routes in hardware.	4120, 5120	64 [q]
	4220	512
	ExtremeSwitching 5520	40,000 [q]
	ExtremeSwitching 5420	6,000
	ExtremeSwitching 5720-MW	107,000 [q]
	7520, 7720	131,000 up to 196,000 [q]
	5720-MXW	213,000 [q]
IPv6 routes with a mask greater than 64 bits in hardware—maximum number of such IPv6 LPM routes in hardware.	5320, 5420	256
	4220, 5520, 7520, 7720	8,192 [r] 32,000 [r]
	5720-MW	16,000 [r]
	5720-MXW	24,000 [r]
IPv6 route sharing in hardware—route mask lengths for which ECMP is supported in hardware.	4120, 4220, 5120, 5320, 5420	0–64, >64 single path only
	5520, 5720, 7520, 7720	0–128 [r]
IP router interfaces—maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	4120, 5120	126
	5320-48T/P, 5420	1,533
	4220, 5320-24T/P, 5320-16P	509
	5320-16P-2MXT-2X	1,021
	5520, 5720, 7520, 7720	2,048
IP multicast static routes—maximum number of permanent multicast IP routes.	All platforms	1,024
IP unicast static routes—maximum number of permanent IP unicast routes.	All platforms	1,024

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
IP route sharing (maximum gateways)—Configurable maximum number of gateways used by equal cost multipath OSPF, BGP, IS-IS, static routes, or L2VPNs. Static routes, OSPF, and BGP are limited to 64 ECMP gateways per destination, while IS-IS is limited to 8. L2VPNs are limited to 16 LSPs per pseudowire on platforms that support 32 gateways, and 64 LSPs per pseudowire on platforms that support 64 gateways.	4120, 4220, 5120, 5320, 5420, 5520	2, 4, or 8
	5720, 7520, 7720	2, 4, 8, 16, 32, or 64

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
IP route sharing (total combinations of gateway sets)—maximum number of combinations of sets of adjacent gateways used by multipath OSPF, BGP, IS-IS, or static routes.	4120, 5120	62 (if maximum gateways is 2, 4, or 8)
	4220, 5320 **Note:** The values here represent the maximum attainable ECMP groups of which, due to the RIOT feature, half are reserved for overlay and half for underlay routing.	124 (if maximum gateways is 2) 124 (if maximum gateways is 4) 60 (if maximum gateways is 8)
	5420 **Note:** The values here represent the maximum attainable ECMP groups of which, due to the RIOT feature, half are reserved for overlay and half for underlay routing.	510 (if maximum gateways is 2) 254 (if maximum gateway is 4) 126 (if maximum gateways is 8)
	5520 **Note:** The values here represent the maximum attainable ECMP groups of which, due to the RIOT feature, half are reserved for overlay and half for underlay routing.	2,046 (if maximum gateways is 2) 1,022 (if maximum gateway is 4) 510 (if maximum gateways is 8)
	5720 if maximum gateways is 2 if maximum gateways is 4 if maximum gateways is 8 if maximum gateways is 16 (default) if maximum gateways is 32 if maximum gateways is 64 **Note:** The values here represent the maximum attainable ECMP groups of which,	2,046 2,046 2,046 1,022 510 254

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
	due to the RIOT feature, half are reserved for overlay and half for underlay routing.	
	7520, 7720 if maximum gateways is 2 if maximum gateways is 4 if maximum gateways is 8 if maximum gateways is 16 (default) if maximum gateways is 32 if maximum gateways is 64 **Note:** The values here represent the maximum attainable ECMP groups of which, due to the RIOT feature, half are reserved for overlay and half for underlay routing.	4,094 4,094 2,046 1,022 510 254
IP multinetting (secondary IP addresses)—maximum number of secondary IP addresses per VLAN.	All platforms	255
Jumbo frames—maximum size supported for jumbo frames, including the CRC.	All platforms	9,216
Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode. **Note:** • The internal lookup table configuration used is "l2-and-l3". • IPv6 and IPv4 L2 IPMC scaling is the same for this mode. • Layer-2 IPMC forwarding cache limits—(IGMP/MLD/PIM snooping) in mixed-mode are the same. 4120 and 4220 do not support PIM snooping.	4120, 5120 4220, 5320 5420 5520 5720-MW 7520, 7720 5720-MXW	16,000 32,000 64,000 32,768 49,152 73,000 81,920

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Layer-3 IPv4 Multicast—maximum number of <S,G,V> entries installed in the hardware (IP multicast compression enabled). **Note:** • Limit value is the same for MVR senders, PIM Snooping entries. PIM SSM cache, IGMP senders, PIM cache. • Assumes source-group-vlan mode as look up key. • Layer 3 IPMC cache limit in mixed mode also has the same value.	4120, 5120	192
	4220	2,000
	5320 (except 5320-24T-4X-XT)	8,000
	ExtremeSwitching 5420M	12,000
	ExtremeSwitching 5420F	6,000
	5520	43,000
	ExtremeSwitching 5720-MW	61,000
	7520, 7720	104,000
	ExtremeSwitching 5720-MXW	110,000
	ExtremeSwitching 5320-24T-4X-XT	2000
Layer-3 IPv6 Multicast—maximum number of <S,G,V> entries installed in the hardware (IP multicast compression enabled). **Note:** • Limit value is the same for MLD sender per switch, PIM IPv6 cache. • Assumes source-group-vlan mode as lookup key. 4120 and 4220 do not support PIM snooping, but MLD cache is supported in the hardware.	4120, 5120	100
	4220	1,000
	ExtremeSwitching 5320 (except 5320-24T-4X-XT)	4,000
	ExtremeSwitching 5420M	6,000
	ExtremeSwitching 5420F	3,000
	ExtremeSwitching 5520	21,500
	ExtremeSwitching 5720-MW	30,500
	7520, 7720	52,000
	ExtremeSwitching 5720-MXW	55,000
	ExtremeSwitching 5320-24T-4X-XT	1,000

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Load sharing—maximum number of load sharing groups. **Note:** The actual number of load-sharing groups that can be configured is limited by the number of physical ports present in the switch or SummitStack.	All platforms	128
Load sharing—maximum number of ports per load-sharing group.	For standalone and stacked: 4120, 4220, 5120, 5320, 5420 For standalone: ExtremeSwitching 5520, 5720, 7520, 7720 For stacked: ExtremeSwitching 5520, 5720, 7520, 7720	8 32 64
Logged messages—maximum number of messages logged locally on the system.	All platforms	20,000
MAC-based security—maximum number of MAC-based security policies.	All platforms	1,024
MAC Locking—Maximum number of MAC locking stations that can be learned on a port.	All platforms	64 (static MAC locking stations) 600 (first arrival MAC locking stations)
Meters—maximum number of meters supported.	All platforms	2,048
Maximum mirroring instances.	All platforms except 4120 and 5120 4120, 5120	4 total, 2 egress 6 defined, max 4 enabled (max 1 egress)
Mirroring (filters)—maximum number of mirroring filters. **Note:** This is the number of filters across all the active mirroring instances.	All platforms	128

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Mirroring, one-to-many (filters)—maximum number of one-to-many mirroring filters. **Note:** This is the number of filters across all the active mirroring instances.	All platforms	128
Mirroring, one-to-many (monitor port)—maximum number of one-to-many monitor ports.	All platforms	16
MLAG ports—maximum number of MLAG ports allowed. **Note:** The number of MLAG ports that can be configured is limited by the number of physical ports present in the system.	5120, 5320	55
	5720	63
	4120, 4220	
	5420, 5520	59
	7520, 7720	61
MLAG peers—maximum number of MLAG peers allowed.	All platforms	2
Multicast listener discovery (MLD) snooping per-VLAN filters—maximum number of VLANs supported in per-VLAN MLD snooping mode.	5320 (except 5320-24T-4X-XT), 5420, 5520, 5720, 7520, 7720	1,500
	4220, 5320-24T-4X-XT	250
	4120, 5120	32
Multicast listener discovery (MLD)v1 subscribers—maximum number of MLDv1 subscribers per port. [n]	5320 (except 5320-24T-4X-XT), 5420, 5520, 5720, 7520, 7720	4,000
	4220, 5320-24T-4X-XT	1,000
	4120, 5120	100
Multicast listener discovery (MLD)v1 subscribers—maximum number of MLDv1 subscribers per switch. [n]	ExtremeSwitching 5320 (except 5320-24T-4X-XT), 5420, 5520	10,000
	ExtremeSwitching 5720-MW	30,000
	7520, 7720	45,000
	ExtremeSwitching 5720-MXW	54,000
	4220, 5320-24T-4X-XT	1,000
	4120, 5120	100
Multicast listener discovery (MLD)v2 subscribers—maximum number of MLDv2 subscribers per port. [n]	ExtremeSwitching 5320 (except 5320-24T-4X-XT), 5420, 5520, 5720, 7520, 7720	4,000
	4220, ExtremeSwitching 5320-24T-4X-XT	1,000
	4120, 5120	100

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Multicast listener discovery (MLD)v2 subscribers—maximum number of MLDv2 subscribers per switch. [n]	4120, 4220, 5320 (except 5320-24T-4X-XT), 5420, 5520	10,000
	ExtremeSwitching 5720-MW	30,000
	7520, 7720	45,000
	ExtremeSwitching 5720-MXW	54,000
	4220, ExtremeSwitching 5320-24T-4X-XT	1,000
	4120, 5120	100
Multicast listener discovery (MLD)v2 maximum source per group—maximum number of source addresses per group.	All platforms except 4120, 5120	200
	4120, 5120	100
Multicast listener discovery (MLD) SSM-map entries—maximum number of MLD SSM mapping entries.	5320, 5420, 5520, 5720, 7520, 7720	500
Multicast listener discovery (MLD) SSM-MAP entries—maximum number of sources per group in MLD SSM mapping entries.	5120, 5320, 5420, 5520, 5720, 7520, 7720	50
Network Address Translation (NAT) VLANs—maximum number of NAT VLANs.	7520, 7720	4
Network Address Translation (NAT) Sessions—number of NAT sessions supported (non twice-NAT).	7520, 7720	1,023
Network Login—maximum number of clients being authenticated on MAC-based VLAN enabled ports.	All platforms	1,024
Network Login—maximum number of dynamic VLANs.	All platforms	1,024
Network Login VLAN VSAs—maximum number of VLANs a client can be authenticated on at any given time.	All platforms	10
Network Service Identifiers (NSI)/VLAN mappings—maximum number of VLANs to NSI mappings.	All platforms	94
Node Alias—maximum number of entries per slot.	All platforms	8,192

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
ONEPolicy Dynamic ACL Rules—maximum number of Dynamic ACLs supported via RADIUS VSA 232 per user in Access-List mode.	All platforms	64
ONEPolicy Roles/Profiles—maximum number of policy roles/profiles.	All platforms	63

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
ONEPolicy Rules per Role/ Profile—maximum number of rules per role/policy.	5320-24T-4X-XT	IPv4 Rules: 256 IPv6 Rules: 0 MAC Rules: 0 L2 Rules: 184
	4120, 5120	IPv4:128 L2:56
	4220	IPv4:256 L2:184
	5320	IPv4 Rules: 1,024 IPv6 Rules: 0 MAC Rules: 0 L2 Rules: 952
	ExtremeSwitching 5420-F, 5320-24T-24S-4XE-XT 7520, 7720	IPv4 Rules: 512 IPv6 Rules: 512 MAC Rules: 512 L2 Rules: 440
	ExtremeSwitching 5720-MW	IPv4 Rules: 1,536 IPv6 Rules: 1,536 MAC Rules: 1,536 L2 Rules: 1,464
	ExtremeSwitching 5720-MXW	IPv4 Rules: 2,048 IPv6 Rules: 2,048 MAC Rules: 2,048 L2 Rules: 1 ,976
	ExtremeSwitching 5420-M, 5520	IPv4 Rules: 1,024 IPv6 Rules: 1,024 MAC Rules: 1,024 L2 Rules: 952

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
ONEPolicy Authenticated Users per Switch—maximum number of authenticated users per switch only with TCI-Overwrite enabled.	ExtremeSwitching 5520, 5720	1,024
	ExtremeSwitching 5320-24T-4X-XT	128
	ExtremeSwitching 5320, 5420, 7520, 7720	512
	4120, 4220, 5120	256
	Stacking	Depends on the stack nodes, but the maximum is 1,024.
ONEPolicy Authenticated Users per Switch—maximum number of authenticated users per switch with TCI-Overwrite disabled. **Note:** The maximum values assume 75% utilization of VLAN-XLATE hash table.	Stacking	1,536–65,534
	7520, 7720	24,576
	ExtremeSwitching 5320-24T-4X-XT	384
	4120, 4220, ExtremeSwitching 5120, 5320, 5420	768
	ExtremeSwitching 5720	12,288
	ExtremeSwitching 5520	9,216
ONEPolicy Authenticated Users per Port per Switch — maximum number of authenticated users per port per switch with TCI overwrite disabled. **Note:** The maximum values assume 75% utilization of VLAN-XLATE hash table.	ExtremeSwitching 5320-24T-4X-XT	384
	4120, 4220, 5120, 5320, 5420	768
	7520, 7720	24,576
	ExtremeSwitching 5720	12,288
	ExtremeSwitching 5520	9,216
ONEPolicy Authenticated Users per Port per Switch — maximum number of authenticated users per port with only with TCI-Overwrite enabled.	4120, 5120	256
	4220	440
	5120, 5320, 5420, 7520, 7720	512
	5520, 5720	1,024
ONEPolicy Permit/Deny Traffic Classification Rules Types—total maximum number of unique permit/deny traffic classification rules types (system/stack).	5320, 5420-F, 7520, 7720	1,976
	5720-MW	6,072
	5720-MXW	8,120
	5420-M, 5520	4,024
	5320-24T-24S-4XE-XT	512
	4220	440
	4120, 5120	184
	5320-24T-4X-XT	128

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
ONEPolicy Permit/Deny Traffic Classification Rules Types—maximum number of unique MAC permit/deny traffic classification rules types (macsource/macdest).	5420-M, 5520	1,024
	5420-F, 5320-24T-24S-4XE-XT 7520, 7720	512
	5720-MW	1,536
	5720-MXW	2,048
	4120, 4220, 5120, 5320	N/A
ONEPolicy Permit/Deny Traffic Classification Rules Types—maximum number of unique IPv6 permit/deny traffic classification rules types (ipv6dest).	ExtremeSwitching 5420-M. 5520	1,024
	ExtremeSwitching 5420-F, 5320-24T-24S-4XE-XT 7520, 7720	512
	ExtremeSwitching 5720-MW	1,536
	ExtremeSwitching 5720-MXW	2,048
	4120, 4220, 5120, 5320	N/A
ONEPolicy Permit/Deny Traffic Classification Rules Types—maximum number of unique IPv4 permit/deny traffic classification rules (typesipsource / ipdest / ipfrag / udpsourceportIP / udpdestportIP / tcpsourceportIP / tcpdestportIP / ipttl / iptos / iptype).	ExtremeSwitching 5320-24T-4X-XT	256
	5120, 5320, 5420-F, 5520	1,024
	ExtremeSwitching 5720-MW	1,536
	ExtremeSwitching 5720-MXW	2,048
	ExtremeSwitching 5420-M, 5320-24T-24S-4XE-XT 7520, 7720	512
	4220	256
	4120, 5120	128
ONEPolicy Permit/Deny Traffic Classification Rules Types—maximum number of unique Layer 2 permit/deny traffic classification rules (ethertype/port).	ExtremeSwitching 5320-24T-24S-4XE-XT	440
	ExtremeSwitching 5320, 5420-M, 5520	952
	5720-MW	1,464
	5720-MXW	1,976
	5420-F, 7520, 7720	440
	4220, 5320-24T-4X-XT	184
	4120, 5120	56
OnePolicy Maximum number of rules supported in AccessList mode—maximum number of rules in AcessList mode.	7520, 7720	3,512
	4120, 5120	440
	4220, 5320-24T-4X-XT	952
	5320, 5420-F, 5320-24T-24S-4XE-XT	4,024
	5420-M	8,120
	5720-MW	12,216
	5720-MXW	16,312

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
OSPFv2/v3 ECMP—maximum number of equal cost multipath OSPFv2 and OSPFv3.	4120, 4220, 5120, 5320, 5420, 5520, 5720	8
	7520, 7720	64
OSPFv2 areas—as an ABR, how many OSPF areas are supported within the same switch.	All platforms	8
OSPFv2 external routes—recommended maximum number of external routes contained in an OSPF LSDB.	5520	80,000
	5720-MW	150,000
	5720-MXW	290,000
	7520, 7720	320,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5420	4,000
	5320-16P-2MXT-2X	992
	4220, 5320-24T-4X-XT	500
	4120, 5120	64
OSPFv2 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	5520	80,000
	5720-MW	150,000
	5720-MXW	290,000
	7520, 7720	320,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5420	1,600
	5320-16P-2MXT-2X	992
	4220, 5320-24T-4X-XT	500
	4120, 5120	64
OSPFv2 combined inter-area and external routes —recommended maximum number of combined inter-area and external routes contained in an OSPF LSDB.	5520	80,000
	5720-MW	150,000
	5720-MXW	290,000
	7520, 7720	320,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5420	1,600
	5320-16P-2MXT-2X	992
	4220, 5320-24T-4X-XT	500
	4120, 5120	64
OSPFv2 inter-vr or leaking routes—recommended maximum number of inter-vr routes contained in an OSPF LSDB.	5420, 5520, 5720, 7520, 7720	2,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT)	1,600
	4120, 5120	64

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
OSPFv2 interfaces—recommended maximum number of OSPF interfaces on a switch (active interfaces only).	All platforms	4
OSPFv2 links—maximum number of links in the router LSA.	4120, 5320, 5420, 5520, 5720, 7520, 7720	400
	4220, 5120	64
OSPFv2 neighbors—maximum number of supported OSPF adjacencies.	All platforms	4
OSPFv2 routers in a single area—recommended maximum number of routers in a single OSPF area.	5520	50
	5720, 7520, 7720	100
	4120, 4220, 5120, 5320, 5420	40
OSPFv2 virtual links—maximum number of supported OSPF virtual links.	All platforms	32
OSPFv3 areas—as an ABR, the maximum number of supported OSPFv3 areas.	5520	16
	5720, 7520, 7720	100
	5120, 5320, 5420	12
OSPFv3 external routes—recommended maximum number of external routes.	5520	40,000
	5720-MW	70,000
	5720-MXW	150,000
	7520, 7720	260,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5720-MW	7,500
	5420	6,000
	5320-24T-4X-XT	300
	5320-16P-2MXT-2X	496
	5120	64

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
OSPFv3 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes.	5520	40,000
	5720-MW	70,000
	5720-MXW	150,000
	7520, 7720	260,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT)	4,000
	5420	6,000
	5320-24T-4X-XT	300
	5320-16P-2MXT-2X	496
	5120	64
OSPFv3 combined inter-area and external routes —recommended maximum number of combined inter-area and external routes.	5520	40,000
	5720-MW	70,000
	5720-MXW	150,000
	7520, 7720	260,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT)	4,000
	5420	6,000
	5320-24T-4X-XT	300
	5320-16P-2MXT-2X	496
	5120	64
OSPFv3 interfaces—maximum number of OSPFv3 interfaces (active interfaces only).	All platforms except 4120 and 4220	4
OSPFv3 neighbors—maximum number of OSPFv3 neighbors.	All platforms except 4120 and 4220	4
OSPFv3 virtual links—maximum number of OSPFv3 virtual links supported.	All platforms except 4120 and 4220	16
PIM IPv4 Limits—maximum number of multicast groups per dynamic rendezvous point.	5120	32
PIM IPv4 Limits—maximum number of multicast groups per static rendezvous point.	All platforms, except 4120 and 5120	180
	4120, 5120	32
PIM IPv4 Limits—maximum number of multicast sources per group.	All platforms except 4120, 4220, 5120	5,000
	4220, 5320-24T-XT	2,000
	4120, 5120	192

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
PIM IPv4 Limits—maximum number of dynamic rendezvous points per multicast group.	All platforms	145
PIM IPv4 Limits—static rendezvous points.	All platforms	32
PIM IPv6 (maximum interfaces)—maximum number of PIM active interfaces.	All platforms	N/A
PIM IPv6 Limits—maximum number of multicast sources per group.	All platforms except 4120, 4220, 5120 4220, 5320-24T-XT 4120, 5120	1,750 1,000 70
PIM IPv6 Limits—maximum number of multicast groups per dynamic rendezvous point.	All platforms except 4120 and 4220	70
PIM IPv6 Limits—maximum number of multicast groups per static rendezvous point.	All platforms except 4120 and 5120 4120, 5120	3,000 (depends on policy file limits) 70
PIM IPv6 Limits—maximum number of dynamic rendezvous points per multicast group.	All platforms	64
PIM IPv6 Limits—maximum number of secondary addresses per interface.	All platforms	70
PIM IPv6 Limits—static rendezvous points.	All platforms	32
Policy-based routing (PBR) redundancy—maximum number of flow-redirects.	All platforms	256 [o]
Policy-based routing (PBR) redundancy—maximum number of next hops per each flow-direct.	All platforms	32 [o]
Port-specific VLAN tags—maximum number of port-specific VLAN tags.	4120, 4220, 5120, 5320, 5420 5520, 5720, 7520, 7720	N/A 1,023
Port-specific VLAN tags—maximum number of port-specific VLAN tag ports.	4120, 4220, 5120, 5320, 5420 5520, 5720, 7520, 7720	N/A 4,000

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Private VLANs—maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN.	4120, 4220, 5120, 5320, 5420, 5520, 5720	36
	7520, 7720	71
Private VLANs—maximum number of private VLANs with an IP address on the network VLAN. **Note:** This limit is dependent on the maximum number of private VLANs in an L2-only environment if the configuration has tagged and translated ports.	4120, 4220, 5120, 5320, 5420, 5520, 5720	960
	7520, 7720	1,024
Private VLANs—maximum number of private VLANs in an L2-only environment.	4120, 4220, 5120, 5320, 5420, 5520, 5720	960
	7520, 7720	1,280
Route policies—suggested maximum number of lines in a route policy file.	All platforms	10,000
RIP Learned Routes—maximum number of RIP routes supported without aggregation.	5320-48T/P, 5320-24T-24S XT, 5420, 5520, 5720, 7520, 7720	10,000
	5320-16P, 5320-24T/P	7,000
	5320-24T-4X-XT	900
	4220, 5320-16P-2MXT-2X	992
	4120, 5120	64
RIP interfaces on a single router—recommended maximum number of RIP routed interfaces on a switch.	All platforms	256
RIPng learned routes—maximum number of RIPng routes.	5320-48T/P, 5320-24T-24S XT, 5420, 5520, 5720, 7520, 7720	3,000
	5120	64
	5320-16P, 5320-24T/P	2,000
	5320-16P-2MXT-2X	496
	5320-24T-4X-XT	400
Spanning Tree (maximum STPDs)—maximum number of Spanning Tree Domains on port mode EMISTP.	5320-48T/P, 5420, 5520, 5720, 5320-24T-24S-4XE-XT, 7520, 7720	64
	4120, 4220, 5120, 5320-24T/P, 5320-16P, 5320-24T-4X-XT	32

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Spanning Tree PVST+—maximum number of port mode PVST domains. **Note:** For all platforms, the maximum number of active ports per PVST domain depends on the maximum number of spanning tree ports supported on given platform. For example, for an ExtremeSwitching switch that supports 256 PVST domains (maximum) and 4,096 STP ports (maximum), the maximum number of active ports per PVST domain would be 16 ports (4,096 ÷ 256).	4120, 4220, 5120, 5320, 5320-24T-4X-XT, 5320-24T-24S-4XE-XT, 5420, 5520, 5720 7520, 7720	128 384
Spanning Tree—maximum number of multiple spanning tree instances (MSTI) domains.	5320-48T/P, 5320-24T-24S-4XE-XT, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P, 5320-24T-4X-XT	64 32
Spanning Tree—maximum number of VLANs per MSTI. **Note:** Maximum number of 10 active ports per VLAN when all 500 VLANs are in one MSTI.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P; 5320-24T-4X-XT, 5320-24T-24S-4XE-XT	600 256
Spanning Tree—maximum number of VLANs on all MSTP instances.	5320-48T/P, 5320-24T-24S-4XE-XT, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P, 5320-24T-4X-XT	1,024 512
Spanning Tree (802.1d domains)—maximum number of 802.1d domains per port.	All platforms	1
Spanning Tree (number of ports)—maximum number of ports including all Spanning Tree domains.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P	4,096 2,048
Spanning Tree (maximum VLANs)—maximum number of STP-protected VLANs (dot1d and dot1w).	5320-48T/P, 5320-24T-24S-4XE-XT, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P, 5320-24T-4X-XT	1,024 600
SSH (number of sessions)—maximum number of simultaneous SSH sessions.	All platforms	8

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
Static MAC multicast FDB entries—maximum number of permanent multicast MAC entries configured into the FDB.	All platforms	1,024
Syslog servers—maximum number of simultaneous Syslog servers that are supported.	All platforms	16
Syslog targets—maximum number of configurable Syslog targets.	All platforms	16
Telnet (number of sessions)—maximum number of simultaneous Telnet sessions.	All platforms	8
Virtual routers—maximum number of user-created virtual routers that can be created on a switch.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P	63 16 (local-only VRs)
Virtual router forwarding (VRFs)—maximum number of VRFs that can be created on a switch. **Note:** * Subject to other system limitations.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P	960 * 16 (local-only VRs)
Virtual router protocols per VR—maximum number of routing protocols per VR.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P	8 N/A
Virtual router protocols per switch—maximum number of VR protocols per switch.	5320-48T/P, 5420, 5520, 5720, 7520, 7720 4120, 4220, 5120, 5320-24T/P, 5320-16P	64 N/A
VLAN aggregation—maximum number of port-VLAN combinations on any one superVLAN and all of its subVLANs.	All platforms	1,000
VLANs—includes all VLANs. **Note:** Only 4,092 user-configurable VLANs are supported. (VLAN 1 is the default VLAN, and 4,095 is the management VLAN, and you may not configure them.)	All platforms	4,094
VLANs (Layer 2)—maximum number of Layer 2 VLANs.	All platforms	4,094

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
VLANs (Layer 3)—maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	5320-48T/P, 5420	1,533
	4120, 5120	126
	4220, 5320-24T/P, 5320-16P	509
	5320-16P-2MXT-2X	1,021
	5520, 5720, 7520, 7720	2,048
VLAN Port Interfaces (VPIF)—maximum number of VLAN port interfaces.	5120, 5320	40,000
	5420	60,000
	4120	32,000
	4220	65,549
	5520, 5720, 7520, 7720	131,585
VLANs (maximum active port-based)—maximum active ports per VLAN when 4,094 VLANs are configured with the default license.	5520, 5720, 7520, 7720	32
	4120, 4220, 5120	15
	5320, 5420	3
VLANs (maximum active protocol-sensitive filters)—number of simultaneously active protocol filters in the switch.	All platforms except 4120 and 4220.	16
VLAN translation—maximum number of translation VLANs. Assumes a minimum of one port per translation and member VLAN.	4120, 4220, 5120, 5320, 5420, 5520, 5720	36
	7520, 7720	71
VLAN translation—maximum number of translation VLAN pairs with an IP address on the translation VLAN. **Note:** This limit is dependent on the maximum number of translation VLAN pairs in an L2-only environment if the configuration includes tagged and translated ports.	4120, 4220, 5120, 5320, 5420, 5520, 5720	960
	7520, 7720	1,024
VLAN translation—maximum number of translation VLAN pairs in an L2-only environment.	4120, 4220, 5120, 5320, 5420, 5520, 5720	960
	7520, 7720	2,046
VMAN CEP—maximum number of CVIDs. **Note:** With 75% hash table utilization.	4120, 4220, 5120, 5320, 5420	768
	5520, 5720	9,000

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
VRRP (v2/v3-IPv4) (maximum instances)—maximum number of VRRP instances for a single switch. **Note:** These limits are applicable for Fabric Routing configuration also. **Note:** Number of groups configured should not exceed the number of individual VRs supported (that is, in normal mode) for that platform type.	**Normal Mode (as individual VRs):**	
	All platforms except 4120, 4220, 5120	511
	4220	508
	4120, 5120	31
	Scaled Mode (with groups):	
	5720, 7520, 7720	2,048
	5120, 5320, 5420, 5520	1,000
	Sliced Mode:	
	All platforms except 4120, 4220, 5120	511
	5120	126
VRRP (v3-IPv6) (maximum instances)—maximum number of VRRP instances for a single switch. (VRRP-VRRPv3-IPv6) **Note:** These limits are applicable for Fabric Routing configuration also. **Note:** Number of groups configured should not exceed the number of individual VRs supported (that is, in normal mode) for that platform type.	**Normal Mode (as individual VRs):**	
	All platforms except 4120, 4220, 5120	511
	4220	508
	4120, 5120	31
	Scaled Mode (with groups):	
	5720, 7520, 7720	2,048
	5120, 5320, 5420, 5520	1,000
VRRP (v2/v3-IPv4/IPv6) (maximum VRID)—maximum number of unique VRID numbers per switch.	All platforms except 4120, 5120	255
	4120, 5120	31
VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN)—maximum number of VRIDs per VLAN.	All platforms except 4120 and 5120	255
	4120, 5120	31
VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks)—maximum number of ping tracks per VLAN.	All platforms	8
VRRP (maximum ping tracks)—maximum number of ping tracks per VRRP Instance under 128 VRRP instances.	All platforms	8 (20 centisecond or 1 second hello interval)

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
VRRP (v3-IPv6) (maximum ping tracks)—maximum number of ping tracks per VRRP Instance under 128 VRRP instances.	All platforms	8 (20 centisecond or 1 second hello interval)
VRRP (v2/v3-IPv4/IPv6) (maximum iproute tracks)—maximum number of IP route tracks per VLAN.	All platforms	8
VRRP (v2/v3-IPv4/IPv6)—maximum number of VLAN tracks per VLAN.	All platforms	8
VXLAN—maximum virtual networks. **Note:** Every VPLS instance/ PSTag VLAN reduces this limit by 1. **Note:** Assumption is all BUM (broadcast/unknown-unicast/multicast) FDB entries are pointing to the same set of RTEPs when all VNETs use explicit flooding. Depends on whether all VNETs use standard or explicit and the number of tenant VLAN ports. **Note:** On ExtremeSwitching 5520 and 5420 switches, every VNET reduces this limit by 1. Every (VPLS/PSTag VLAN) + port reduces the limit by 1 on all platforms. Every VXLAN Underlay Multicast Tunnel reduces this limit by 1.	5520, 5720, 7520, 7720 5320, 5420	2,048–4,000 150-375
VXLAN—maximum tenant VLANs plus port combinations **Note:** Every (VPLS/PSTag VLAN) + port reduces the limit by 1.	5520, 5720, 7520, 7720 5320, 5420	4,096 150-375
VXLAN—maximum static MAC to IP bindings. **Note:** Every FDB entry configured reduces this limit by 1.	All supported platforms except 4120 and 4220	64,000

Table 7: Supported Limits for the Base License (continued)

Metric	Product	Limit
VXLAN—maximum RTEP IP addresses	All supported platforms except 4120 and 4220	512
VXLAN—maximum virtual networks with dynamic learning and OSPF extensions for VXLAN	5520, 5720, 7520, 7720 5320, 5420	4,000 375
VXLAN—or replicator role, maximum number of attached leafs per switch.	All supported platforms except 4120 and 4220	256
XML requests—maximum number of XML requests per second. **Note:** Limits are dependent on load and type of XML request. These values are dynamic ACL data requests.	All platforms	10 with 100 DACLs
XNV authentication—maximum number of VMs that can be processed (combination of local and network VMs).	All platforms except 4120 and 4220	2,048
XNV database entries—maximum number of VM database entries (combination of local and network VMs).	All platforms except 4120 and 4220	16,000
XNV database entries—maximum number of VPP database entries (combination of local and network VPPs).	All platforms except 4120 and 4220	2,048
XNV dynamic VLAN—Maximum number of dynamic VLANs created (from VPPs /local VMs).	All platforms except 4120 and 4220	2,048
XNV local VPPs—maximum number of XNV local VPPs.	All platforms except 4120 and 4220	2,048 ingress 512 egress
XNV policies/dynamic ACLs—maximum number of policies/ dynamic ACLs that can be configured per VPP.	All platforms except 4120 and 4220	8 ingress 4 egress
XNV network VPPs—maximum number of XNV network VPPs.p	All platforms except 4120 and 4220	2,048 ingress 512 egress

Premier License Limits

The following table shows supported limits for features in the Premier License.

Table 8: Supported Limits for the Premier License

Metric	Product	Limit
Anycast RP Using PIM—maximum number of IPv4 Anycast RP set per VR.	All platforms	32
Anycast RP Using PIM—maximum number of IPv6 Anycast RP set per VR.	All platforms	32
Anycast RP Using PIM—RP peers per Anycast RP set.	All platforms	10
BGP (aggregates)—maximum number of BGP aggregates.	5120, 5320, 5420, 5520, 5720, 7520, 7720	256
BGP (networks)—maximum number of BGP networks.	5120, 5320, 5420, 5520, 5720, 7520, 7720	1,024
BGP (peers)—maximum number of BGP peers. **Note:** With default keepalive and hold timers. **Note:** Each BGPv4/BGPv6 peer handles a maximum of 50 routes. **Note:** ECMP should not be enabled for BGP.	5120, 5320, 5420, 5520, 5720, 7520, 7720	300
BGP (peer groups)—maximum number of BGP peer groups.	5120, 5320, 5420, 5520, 5720, 7520, 7720	64
BGP (policy entries)—maximum number of BGP policy entries per route policy.	5120, 5320, 5420, 5520, 5720, 7520, 7720	256
BGP (policy statements)—maximum number of BGP policy statements per route policy.	5120, 5420, 5520, 5720, 7520, 7720 5320	1,024 820

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
BGP (multicast address-family routes)—maximum number of multicast address-family routes.	5520, 5720MW	13,000
	5720-MXW	20,000
	7520, 7720	25,000
	5320 48-port, 5420	12,000
	5320-16P-4XE, 5320 24-port except XT	8,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	992
	5120	64
BGP (unicast address-family routes)—maximum number of unicast address-family routes.	5520, 5720MW (at default)	13,000
	7520, 7720 (at default)	25,000
	5720-MXW (at default)	20,000
	5320 48-port, 5420	12,000
	5320-16P-4XE, 5320 24-port except XT	8,000
	5120	64
	5320-24T-4X-XT, 5320-16P-2MXT-2X	992
	5720-MW (with ALPM enabled)	163,000
	5720-MXW (with ALPM enabled)	288,000
	5520 (with ALPM enabled)	80,000
BGP (non-unique routes)—maximum number of non-unique BGP routes.	7520, 7720	75,000
	5320 48-port, 5420, 5520, 5720-MW	36,000
	5720-MXW	60,000
	5320-16P-4XE, 5320 24-port except XT	24,000
	5120	192
	5320-24T-4X-XT, 5320-16P-2MXT-2X	2,700
BGP ECMP—maximum number of equal cost paths per multipath for BGP and BGPv6.	5120, 5320, 5420, 5520, 7520, 7720	8
	5720	64

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
BGPv6 (unicast address-family routes)—maximum number of unicast address family routes.	5320 48-port, 5420, 5520, 5720-MW (at default)	6,000
	5720-MW (with ALPM enabled)	107,000
	5720-MXW, 7520, 7720 (at default)	10,000
	5720-MXW (with ALPM enabled)	213,000
	5320-16P-4XE, 5320 24-port except XT	4,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	496
	5120	64
	5520 (with ALPM enabled)	40,000
BGPv6 (non-unique routes)—maximum number of non-unique BGP routes.	5420, 5520, 5720-MW	18,000
	5720-MXW, 7520, 7720	30,000
	5320 (except 5320-24T-4X-XT, 5320-16P-2MXT-2X)	14,000
	5320-24T-4X-XT, 5320-16P-2MXT-2X	1,488
	5120	192
EVPN EVI instances—maximum number of EVI instances.	All platforms, except 4120 and 5120	1,024
IS-IS adjacencies—maximum number of supported IS-IS adjacencies.	All platforms	128
IS-IS ECMP—maximum number of equal cost paths per multipath for IS-IS.	All platforms	2, 4, or 8
IS-IS interfaces—maximum number of interfaces that can support IS-IS.	All platforms	255
IS-IS routers in an area—recommended maximum number of IS-IS routers in an area.	All platforms	256
IS-IS route origination—recommended maximum number of routes that can be originated by an IS-IS node.	All platforms	20,000
IS-IS IPv4 L1 routes in an L1 router—recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	All platforms	25,000

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
IS-IS IPv4 L2 routes—recommended maximum number of IS-IS Level 2 routes.	All platforms	25,000
IS-IS IPv4 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in an L1/L2 IS-IS router.	All platforms	20,000
IS-IS IPv6 L1 routes in an L1 router—recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	All platforms	10,000
IS-IS IPv6 L2 routes—recommended maximum number of IS-IS Level 2 routes.	All platforms	10,000
IS-IS IPv6 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in a L1/I2 router.	All platforms	10,000
IS-IS IPv4/IPv6 L1 routes in an L1 router—recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	All platforms	20,000
IS-IS IPv4/IPv6 L2 routes in an L2 router—recommended maximum number of IS-IS Level 2 routes in a Level 2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	All platforms	20,000
IS-IS IPv4/IPv6 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in a Level 1/Level2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	All platforms	20,000

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
L2 VPN: VCCV (pseudowire Virtual Circuit Connectivity Verification) VPNs per switch—maximum number of VCCV enabled VPLS VPNs.	5520, 7520, 7720 5120, 5320, 5420, 5720	16 N/A
L2 VPN: VPLS MAC addresses—maximum number of MAC addresses learned by a switch.	5520 7520, 7720 5120, 5320, 5420, 5720	64,000 140,000 N/A
L2 VPN: VPLS VPNs—maximum number of VPLS virtual private networks per switch.	5520, 7520, 7720 5120, 5320, 5420, 5720	1,023 N/A
L2 VPN: VPLS peers—maximum number of VPLS peers per VPLS instance.	5520, 7520, 7720 5120, 5320, 5420, 5720	64 N/A
L2 VPN: LDP pseudowires—maximum number of pseudowires per switch.	5520 7520, 7720 5120, 5320, 5420, 5720	3,500 7,000 N/A
L2 VPN: static pseudowires—maximum number of static pseudowires per switch.	5520 7520, 7720 5120, 5320, 5420, 5720	3,500 7,000 N/A
L2 VPN: Virtual Private Wire Service (VPWS) VPNs—maximum number of virtual private networks per switch.	5520 7520, 7720 5120, 5320, 5420, 5720	1,023 4,090 N/A
MPLS RSVP-TE interfaces—maximum number of interfaces.	5520, 7520, 7720 5120, 5320, 5420, 5720	32 N/A
MPLS RSVP-TE ingress LSPs—maximum number of ingress LSPs.	5520, 7520, 7720 5120, 5320, 5420, 5720	2,000 N/A
MPLS RSVP-TE egress LSPs—maximum number of egress LSPs.	5520, 7520, 7720 5120, 5320, 5420, 5720	2,000 N/A
MPLS RSVP-TE transit LSPs—maximum number of transit LSPs.	5520, 7520, 7720 5120, 5320, 5420, 5720	4,000 N/A
MPLS RSVP-TE paths—maximum number of paths.	5520 7520, 7720 5120, 5320, 5420, 5720	1,000 2,000 N/A

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
MPLS RSVP-TE profiles—maximum number of profiles.	5520	1,000
	7520, 7720	2,000
	5120, 5320, 5420, 5720	N/A
MPLS RSVP-TE EROs—maximum number of EROs per path.	5520, 7520, 7720	64
	5120, 5320, 5420, 5720	N/A
MPLS LDP peers—maximum number of MPLS LDP peers per switch.	5520, 7520, 7720	128
	5120, 5320, 5420, 5720	N/A
MPLS LDP adjacencies—maximum number of MPLS LDP adjacencies per switch.	5520, 7520, 7720	64
	5120, 5320, 5420, 5720	N/A
MPLS LDP ingress LSPs—maximum number of MPLS LSPs that can originate from a switch.	5520, 7520, 7720	2,048
	5120, 5320, 5420, 5720	N/A
MPLS LDP-enabled interfaces—maximum number of MPLS LDP configured interfaces per switch.	5520, 7520, 7720	128
	5120, 5320, 5420, 5720	N/A
MPLS LDP transit LSPs—maximum number of MPLS transit LSPs per switch.	5520	3,500
	7520, 7720	4,000
	5120, 5320, 5420, 5720	N/A
MPLS LDP egress LSPs—maximum number of MPLS egress LSPs that can terminate on a switch.	5520	3,500
	7520, 7720	4,000
	5120, 5320, 5420, 5720	N/A
MPLS static egress LSPs—maximum number of static egress LSPs.	5520	3,500
	7520, 7720	8,000
	5120, 5320, 5420, 5720	N/A
MPLS static ingress LSPs—maximum number of static ingress LSPs.	5520	3,500
	7520, 7720	4,000
	5120, 5320, 5420, 5720	N/A
MPLS static transit LSPs—maximum number of static transit LSPs	5520	3,500
	7520, 7720	4,000
	5120, 5320, 5420, 5720	N/A
MSDP active peers—maximum number of active MSDP peers.	5120	16
	5320, 5420, 5520, 5720, 7520, 7720	64

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
MSDP SA cache entries—maximum number of entries in SA cache.	5120	192
	5320, 5420F	6,000
	5420M	8,000
	5520, 5720, 7520, 7720	14,000
MSDP maximum mesh groups—maximum number of MSDP mesh groups.	All platforms	16
OSPFv2/v3 ECMP—maximum number of equal cost multipath OSPFv2 and OSPFv3.	5120, 5320, 5420, 5520	8
	5720	64
OSPFv2 areas—as an ABR, how many OSPF areas are supported within the same switch.	All platforms	8
OSPFv2 external routes—recommended maximum number of external routes contained in an OSPF LSDB.	5120	64
	5520	80,000
	5720-MX	150,000
	5720-MXW	290,000
	7520, 7720	320,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5420	4,000
	5320-16P-2MXT-2X	992
	5320-24T-4X-XT	500
OSPFv2 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	5120	64
	5520	80,000
	5720-MW	150,000
	5720-MXW	290,000
	7520, 7720	320,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5420	1,600
	5320-16P-2MXT-2X	992
	5320-24T-4X-XT	500

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
OSPFv2 combined inter-area and external routes —recommended maximum number of combined inter-area and external routes contained in an OSPF LSDB.	5120	64
	5520	80,000
	5720-MW	150,000
	5720-MXW	290,000
	7520, 7720	320,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5420	1,600
	5320-16P-2MXT-2X	992
	5320-24T-4X-XT	500
OSPFv2 inter-vr or leaking routes—recommended maximum number of inter-vr routes contained in an OSPF LSDB.	5420, 5520, 5720, 7520, 7720	2,000
	5320 (5320-16P-2MXT-2X, 5320-24T-4X-XT)	1,600
	5120	64
OSPFv2 interfaces—recommended maximum number of OSPF interfaces on a switch (active interfaces only).	5420, 5520, 5720, 7520, 7720	400
	5320	320
	5120	64
OSPFv2 links—maximum number of links in the router LSA.	5420, 5520, 5720, 7520, 7720	400
	5320	320
	5120	64
OSPFv2 neighbors—maximum number of supported OSPF adjacencies.	5420, 5520, 5720, 7520, 7720	128
	5320	96
	5120	64
OSPFv2 routers in a single area—recommended maximum number of routers in a single OSPF area.	5420, 5520	50
	5720, 7520, 7720	100
	5120, 5320	40
OSPFv2 virtual links—maximum number of supported OSPF virtual links.	5420, 5520, 5720, 7520, 7720	32
	5120, 5320	25
OSPFv3 areas—as an ABR, the maximum number of supported OSPFv3 areas.	5420, 5520	16
	5720, 7520, 7720	100
	5120, 5320	12

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
OSPFv3 external routes—recommended maximum number of external routes.	5520	40,000
	5720-MW	70,000
	5720-MXW	150,000
	7520, 7720	260,000
	5120, 5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5720-MW	7,500
	5420	6,000
	5320-16P-2MXT-2X	496
	5320-24T-4X-XT	300
	5120	64
OSPFv3 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes.	5520	40,000
	5720-MW	70,000
	5720-MXW	150,000
	7520, 7720	260,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT)	4,000
	5420	6,000
	5320-16P-2MXT-2X	496
	5320-24T-4X-XT	300
	5120	64
OSPFv3 combined inter-area and external routes—recommended maximum number of combined inter-area and external routes.	5520	40,000
	5720-MW	70,000
	5720-MXW	150,000
	7520, 7720	260,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT)	4,000
	5420	6,000
	5320-16P-2MXT-2X	496
	5320-24T-4X-XT	300
	5120	64
OSPFv3 interfaces—maximum number of OSPFv3 interfaces (active interfaces only).	5420, 5520, 5720, 7520, 7720	256
	5320	192
	5120	64
OSPFv3 neighbors—maximum number of OSPFv3 neighbors.	5420, 5520, 5720, 7520, 7720	64
	5120, 5320	48

Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
OSPFv3 virtual links—maximum number of OSPFv3 virtual links supported.	All platforms except 4120 and 4220	16
PIM IPv4 (maximum interfaces)—maximum number of PIM active interfaces.	5320, 5420, 5520, 5720, 7520, 7720	255
	5120	60
PIM IPv4 Limits—maximum number of multicast groups per dynamic rendezvous point.	5120, 5320, 5420, 5520, 5720, 7520, 7720	180
PIM IPv4 Limits—maximum number of multicast groups per static rendezvous point.	5320, 5420, 5520, 5720, 7520, 7720	3,000 (depends on policy file limits)
	5120	192
PIM IPv4 Limits—maximum number of multicast sources per group.	5320, 5420, 5520, 5720, 7520, 7720	5,000
	5120	192
PIM IPv4 Limits—maximum number of dynamic rendezvous points per multicast group.	5320, 5420, 5520, 5720, 7520, 7720	145
	5120	32
PIM IPv4 Limits—static rendezvous points.	5120, 5320, 5420, 5520, 5720, 7520, 7720	32
PIM IPv6 (maximum interfaces)—maximum number of PIM active interfaces.	5320, 5420, 5520, 5720, 7520, 7720	255
	5120	30
PIM IPv6 limits—maximum number of multicast sources per group.	5320, 5420, 5520, 5720, 7520, 7720	1,750
	5120	70
PIM IPv6 limits—maximum number of multicast groups per dynamic rendezvous point.	5120, 5320, 5420, 5520, 5720, 7520, 7720	70
PIM IPv6 limits—maximum number of multicast groups per static rendezvous point.	5320, 5420, 5520, 5720, 7520, 7720	3,000 (depends on policy file limits)
	5120	70
PIM IPv6 limits—maximum number of multicast groups per dynamic rendezvous points per multicast group.	5320, 5420, 5520, 5720, 7520, 7720	64
	5120	20

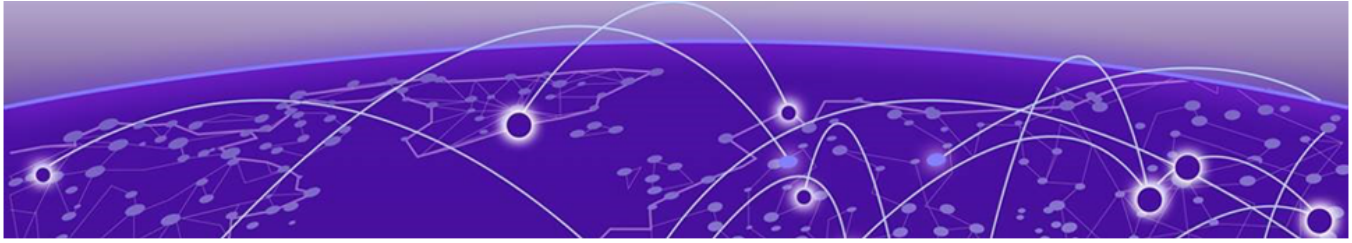
Table 8: Supported Limits for the Premier License (continued)

Metric	Product	Limit
PIM IPv6 limits—maximum number of secondary addresses per interface	5320, 5420, 5520, 5720, 7520, 7720 5120	70 30
PIM IPv6 limits—maximum number of static rendezvous points.	5120, 5320, 5420, 5520, 5720, 7520, 7720	32
PTP/1588v2 Clock Ports	7520-48Y, 7720-32C	32 for boundary clock
PTP/1588v2 Clock Instances	5320-16P-2MXT, 5420, 5520, 5720, 7520-48Y, 7720 7520-48Y, 7720-32C	1 transparent clock 1 boundary clock
PTP/1588v2 Unicast Static Masters	7520-48Y, 7720-32C	10 entries per Slave clock type

Notes for Limits Tables

[a] The table shows the total available. When installing ACL rules bound to a set of ports, rules are replicated for each port if there are ACL counters and counter compression is not enabled, or if the ports are extended ports.
[c] When there are BFD sessions with minimal timer, sessions with default timer should not be used.
[f] Effective capacity varies based on actual MAC addresses and VLAN IDs used and hash algorithm selected.
[g] Based on "configure forwarding internal-tables more l2".
[h] Based on "configure forwarding internal-tables more l3-and-ipmc".
[j] The limit depends on setting configured with configure iproute reserved-entries.
[m] The IPv4 and IPv6 multicast entries share the same hardware tables, so the effective number of IPv6 multicast entries depends on the number of IPv4 multicast entries present and vice versa.
[n] If IGMP and MLD are simultaneously configured on the switch, the number of effective subscribers supported are lessened accordingly.
[o] The total of all PBR next hops on all flow redirects should not exceed 4,096.

[p] The number of XNV authentications supported based on system ACL limitations.
[q] Based on "configure forwarding internal-tables more routes".
[r] Based on `configure forwarding internal-tables more routes ipv6-mask-length 128`.
[s] Based on `configure forwarding internal-tables more l3-and-ipmc` or `configure forwarding internal-tables l2-and-l3`.



Open Issues, Known Behaviors, and Resolved Issues

Open Issues on page 105
Known Behaviors on page 106
Resolved Issues in Switch Engine on page 107

This topic lists open software issues, limitations in system architecture (known issues), and resolved issues in Switch Engine.

Open Issues

The following are new open issues for supported features found in this version.

Table 9: Open issues in version 34.1.1

Defect Number	Description
General	
CFD-15428	AVB streams are randomly dropped on 5520-24X switches.
CFD-17326	Packet loss is seen when MLAG switch with reload-delay configuration is rebooted.
CFD-17445	On a 7520-48XT, the stack Ports aren't programmed properly when using a mix of "native-v400" and "native-v400 alternative-configuration".
CFD-17488	Switch's boot menu displays same image name in both partitions although the installed images are different.
CFD-18011	"IPv4mc Grp table full" error was seen even though the number of programmed entries had not reached the table limit.
CFD-18391	EXOS secure syslog over TLS/TCP 6514 intermittently stops forwarding until connection reset.
CFD-18423	FEC operational status set to CL91 after upgrade to 33.7.1.6 despite FEC configured off.
CFD-18458	FA Authentication configuration occurs for Mgmt port which is not supported.

Table 9: Open issues in version 34.1.1 (continued)

Defect Number	Description
CFD-18504	While running cable diagnostics for a single port using the command `run diagnostics cable port` *`port_number`*, it initiates cable diagnostics for all ports of the switch.
CFD-18643	EAPS instability when using "open-secondary-port" with MACsec.
CFD-18657	The **`watch`** command buffers outputs of the runs are complete and display it in a single burst, instead of displaying it after each execution.
CFD-18661	Port utilization counters showing erroneous/ impossible values.
CFD-18823	On 5120 switches, "Failed to load configuration (ESRP is still not ready yet)" error message is occasionally logged after the switch boots up.

Known Behaviors

The following table lists limitations in system architecture that have yet to be resolved.

Table 10: Known Issues, Platform-Specific, and Feature Change Requests (CRs) in 34.1.1

Defect Number	Description
EXOS-39972	IPv4 Source IP Lockdown Blocks DHCPv6 Snooping Binding Learning: If IPv4 source IP lockdown is enabled on a port and IPv6 source IP lockdown is not, DHCPv6 snooping on that port does not learn IPv6 bindings for the clients on the port. IPv4 and IPv6 source IP lockdown share one deny rule on each port, so the rule drops the IPv6 Neighbor Discovery traffic that the DHCPv6 server needs to resolve the client. The server therefore never sends the Advertise or Reply message. This limitation occurs in a Layer 2 setup (DHCPv6 server, switch, client). It does not occur in a DHCPv6 relay setup. Workaround: When you use DHCPv6 snooping, also enable IPv6 source IP lockdown on the same ports: `enable ip-security ipv6 source-ip-lockdown ports` *`port_list`* In a Layer 2 setup, do not enable source IP lockdown on the port that faces the DHCPv6 server.

Resolved Issues in Switch Engine

The following issues were resolved in version 34.1.1. Version 34.1.1 includes all fixes up to and including versions 31.6, 31.7, 32.1, 32.2, 32.3, 32.4, 32.5, 32.6.x, 32.7.x, 33.2.1, 33.3.1, 33.4.1, 33.5.x, 33.6.x, and 33.7.x.

Table 11: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in 34.1.1

Defect Number	Description
General	
CFD-14144	The **`show switch management`** command displays image integrity checking as unknown instead of valid.
CFD-14421	FDB was not learned from the received IGMP report when the switch is in SW learning mode.
CFD-14863	Broadcast and Multicast counters are not incremented in the **`show port mgmt stat`** command.
CFD-15771	Need provision for "Port-Based FDB-Flush" in EAPS implementation, which can be used in scaled networks to minimize hardware programming load.
CFD-15797	SNMP Error was seen after upgrade to version 32.7.3.
CFD-15926	The **`restart auto-provision cloud-connector`** command needs to be introduced, which will be used to by ExtremeCloud™ IQ Site Engine for license refresh.
CFD-16425	A port disabled by ELRP permanently due to a loop in a dynamic VLAN is re-enabled as soon as the dynamic VLAN is deleted.
CFD-16501	When one of the stack member is removed/rebooted, **`show fans`** displays "operational at 0 RPM" instead of Empty
CFD-16672	SFLOW sub-sampling is set to 2 even though all ports are configured with the same sample-rate
CFD-16675	Improved log message when an Access profile denies access based on an implicit deny without a deny entry in dynamic ACL.
CFD-16874	PIM join messages was not processed by the RP when it receives (S, G, RPT) after the interface pruned.
CFD-16898	The MRP process crashes when simultaneously enabling MVRP on an MLAG port on both MLAG peers.
CFD-16939	BGP blackhole entry remains in the kernel after removing the corresponding NLRI entry from the BGP route policy.
CFD-16988	Rx Jabber is incrementing even when port is disabled.
CFD-17014	On switches that don't have a dedicated OOB mgmt port, a memory leak occurs in the nettools process whenever cloud-connector is run.
CFD-17056	NTP configuration was lost after reboot if DHCP is enabled on the VLAN.
CFD-17172	VXLAN traffic not passing through a CEP CVID port of a VMAN.
CFD-17311	cfgmgr process crash noticed during license check on Universal platforms if the switch has user-created VR in the configuration.

Table 11: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in 34.1.1 (continued)

Defect Number	Description
CFD-17407	Help-text of **debug hal show optic-info** needs to be updated.
CFD-17522	VLAN process crashes due to exdb infrastructure handler.
CFD-17872	The command **show switch non-volatile-storage health-check** is restricted for non-admin users.
CFD-17925	Password for user admin is not getting updated when trying to change it using Web GUI / Chalet
CFD-17994	Timezone configuration in the `show configuration devmgr` command output displays autodst even when it is not enabled explicitly.
CFD-18200	When the switch is in software-learning mode and with netlogin enabled, IGMP report packet with "Router-Alert" option are ignored.
CFD-18345	Netlogin memory leak was seen due to authRequest.
CFD-18356	Process cliMaster crashes when **restart process exsshd** is issued in an SSH session, leading to switch resets.
CFD-18422	POE process crash was seen in X440-G2.
CFD-18464	The OSPF process crashes when the **show ospf lsdb detail** command is executed simultaneously on multiple CLI sessions.
EXOS-39816	In MLAG environment, packet blocking is not functioning correctly, resulting in duplicate packet forwarding and Layer 2 loops.
EXOS-39930	In ExtremeCloud™ IQ, collection for device not complete after adding SSH access profiles in the switch.
VPEX	
CFD-16520	When a 5420 is a controlling bridge, forwarding issues might occur on certain ports if more than 500 BPE ports are added as untagged to VLANs.
4220 Series Switches	
CFD-15685	Telemetry rules not installed on re-configuration.
5320 Series Switches	
CFD-16392	Extended Diagnostics fails when a link is up at 100M.
5520 Series Switches	
CFD-14188	VXLAN packets are being processed by the CPU.
CFD-16177	OnePolicy does not filter incoming IGMP on 5520.
7520 Series Switches	
CFD-18328	25G ports intermittently fail to come up on 7520 switches.