Switch Engine v33.7.1-Patch1-26 Release Notes

New Features, Improvements, and Known Issues

9041140-02 Rev AA
September 2026

Legal Notice

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

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

Trademarks

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

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

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

Open Source Declarations

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

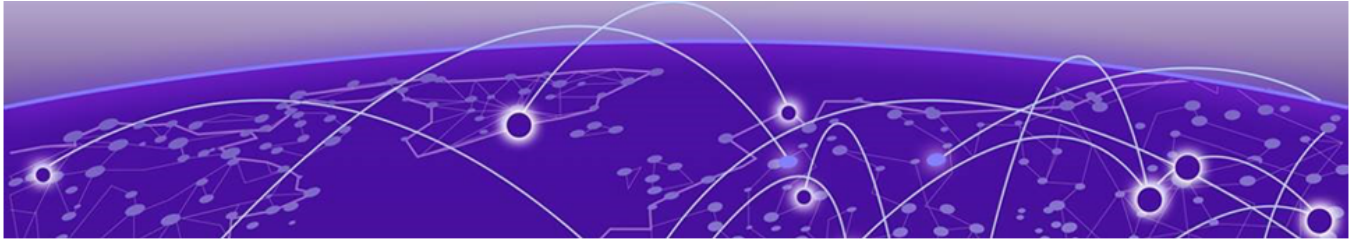

Table of Contents

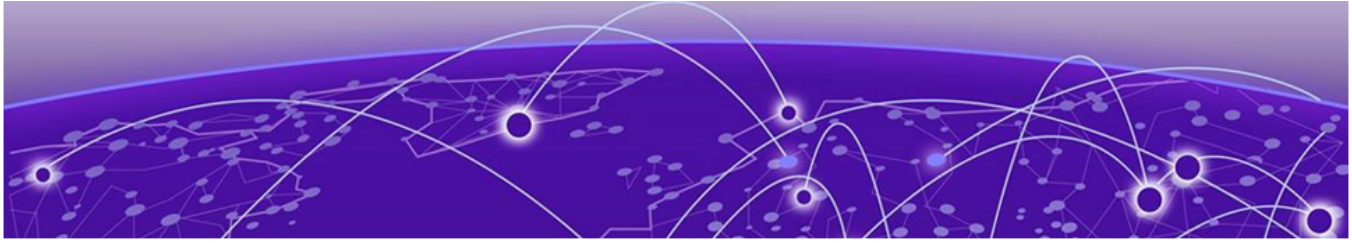

Abstract

Switch Engine v33.7.1-Patch1-26 Release Notes by Extreme Networks, Inc., released in September 2026, provide comprehensive details on new features, software improvements, scaling limits, and resolved issues in version 33.7.1.6-Patch1-26. Key technical enhancements include support for configuring an alternate MAC address, improvements in Fabric Attach conditional management VLAN creation and server synchronization, new capabilities for ACL packet log rate limiting at both global and rule levels, configurable burst tolerance and out-of-profile actions for flood rate-limit control, a new command to optimize EAPS (Ethernet Automatic Protection Switching) convergence performance in large-scale deployments by providing configurable FDB flush methods, and enhanced ExtremeCloud IQ Agent logging that defaults to in-memory buffer-based storage with on-demand file-based logging. Additional improvements include new CLI commands for various functionalities, enhancements in Fabric Attach timeout settings, and refined diagnostic capabilities. The release notes outline hardware and software compatibility, default settings, and image file names, along with guidance for upgrading Switch Engine; limits for various licenses and features, including Base and Premier licenses, are detailed. The notes highlight known behaviors and limitations in the system architecture, and list numerous resolved issues across different patches, including improvements in security profile operation and protocol handling. This release serves as a comprehensive resource for technical readers seeking detailed insights into the functionality, compatibility, and performance improvements of the specified software version."

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.

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.

Help and Support

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

Extreme Portal

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

The Hub

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

Call GTAC

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

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

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

Subscribe to Product Announcements

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

1. Go to The Hub.
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.

4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

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

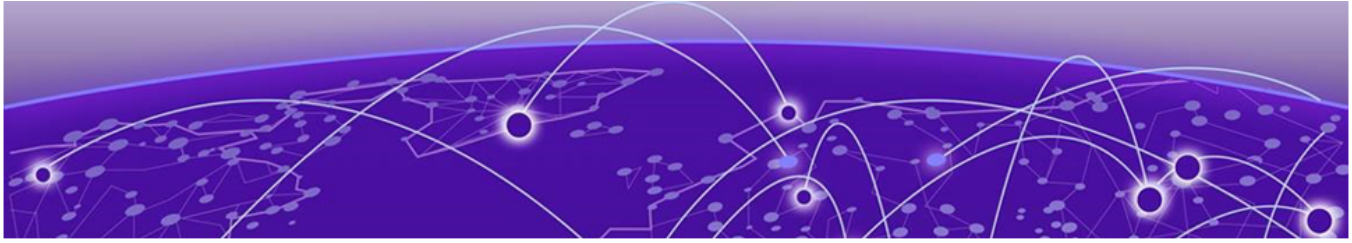

Overview

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

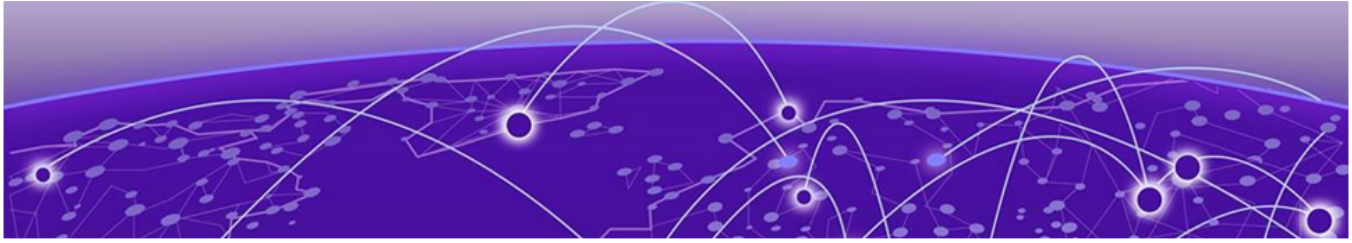

Security Information

The following section covers important security information.

Linux Kernel

This version of Switch Engine uses Linux Kernel 5.10.

OpenSSL Version

This version of Switch Engine uses FIPS openssl-3.0.10.

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
```

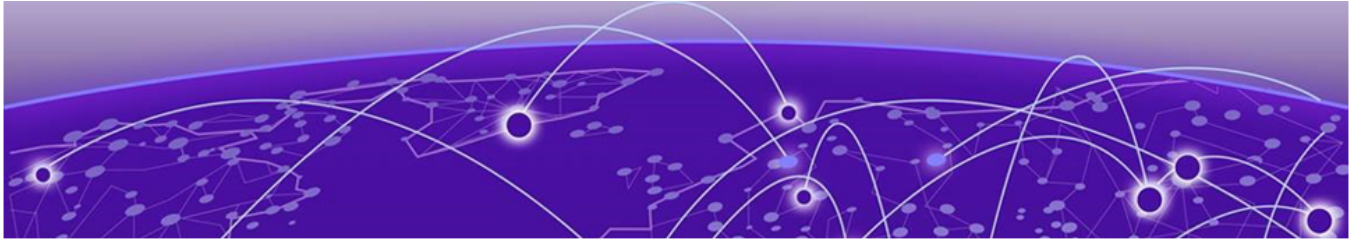

Newly Purchased Switches Require Software Upgrade

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.

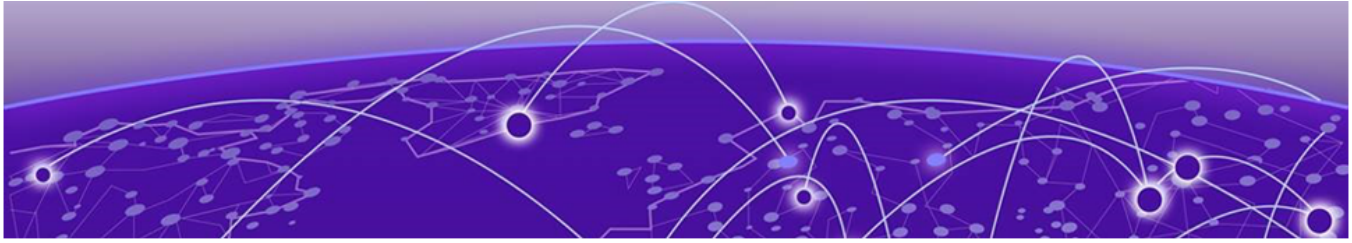

Default Settings

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]	

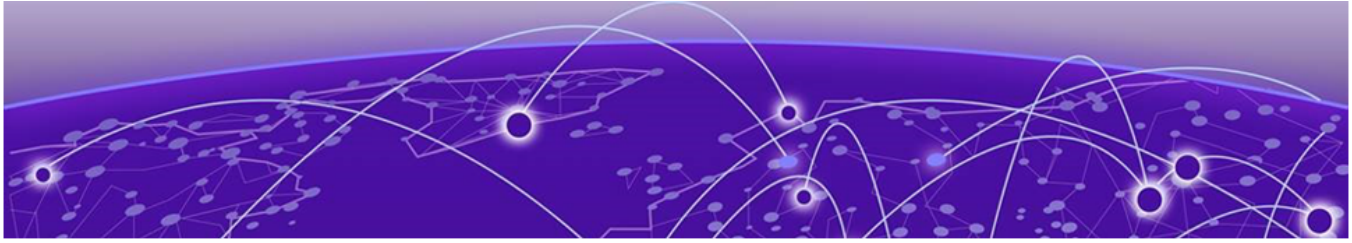

Switch Engine Image File Names

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`

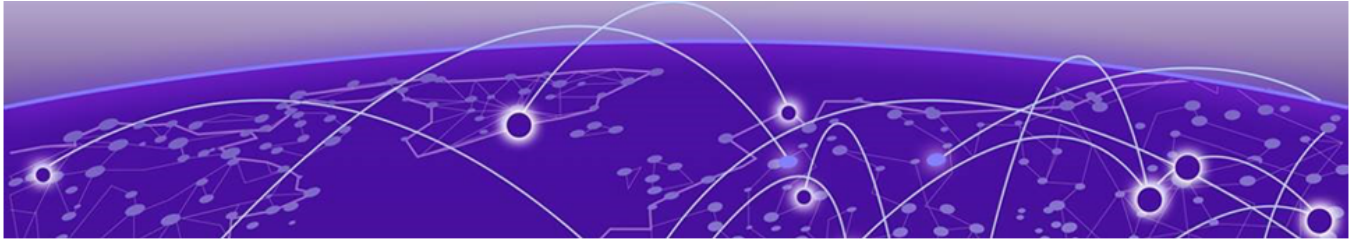

New and Corrected Features

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

5120 Series MACsec Support

Version 33.7.1 introduces MACsec (IEEE 802.1AE) support on 5120 Series switches, enabling secure Layer 2 encryption for Ethernet traffic. This feature provides hardware-based encryption and integrity protection for point-to-point links, helping secure data in transit across trusted and untrusted domains.

MACsec support includes:

- Secure connectivity using AES-GCM encryption
- Support for MKA (MACsec Key Agreement) protocol
- Integration with 802.1X authentication workflows

Verify platform-specific limitations and licensing requirements before deployment.

ACL Packet Log Rate Limiting

The ACL packet log rate-limiting feature provides configurable control over the logging frequency for packets that match ACL rules with `log` and `log-raw` action modifiers. Configure rate limits at both global and individual rule levels to prevent log saturation and manage logging facility resources.

Issue: Packet logs generated by ACL rule action modifiers `log` and `log-raw` have a fixed rate limit of 10 logs per 5 seconds. This fixed rate prevents operators from adjusting logging

intensity based on individual rules or network requirements. A single high-frequency rule can monopolize the EMS and Syslog logging facilities, preventing logs from other rules and system events from being recorded.

Resolution: The ACL packet log rate-limiting feature allows configuration at two levels that work together: a global default that limits all ACL packet logs, and rule-specific rate limits that prevent individual rules from consuming excessive logging facility resources. Operators can specify burst size and time interval parameters at each level to control logging frequency based on deployment requirements and traffic patterns.

Key Capabilities:

- Global rate-limit configuration sets the maximum allowable rate for all ACL packet logs from rules using `log` and `log-raw` action modifiers
- Rule-level rate-limit configuration limits the logging frequency for individual ACL entries, preventing a single high-frequency rule from monopolizing the EMS and Syslog logging facilities
- Rule-level rate limiters feed into the global rate limiter; logs from all rules are subject to the global rate limit regardless of rule-level configuration
- Configurable burst size and time interval parameters allow precise control over logging frequency at each level
- Rules default to the global rate-limit setting when no rule-level configuration is specified

Configuration Syntax

Global Rate Limiting

Configure the default rate-limit policy for all ACL packet logging:

```
configure access-list action-modifier log rate-limit <burst_size> in <time_interval>
```

Rule-Level Rate Limiting

Configure the rate-limit for a specific ACL rule by including the rate-limit parameter in the rule action block:

```
log { rate-limit <burst_size> in <time_interval> }
```

Example Rule Configuration

```
entry packet_log { if { destination-address 10.1.1.0/24; protocol udp; } then { log rate-limit 1 in 5; mirror-cpu; } }
```

This rule logs one packet per 5 seconds for UDP traffic destined to 10.1.1.0/24. Logs from this rule and all other rules are still subject to the global rate-limit setting.

Supported Platforms

All platforms.

Automate RADIUS Reauthentication Pause Using Passive Server Reachability Detection

Version 33.7.1 adds automated, passive detection of RADIUS server availability to trigger and clear the RADIUS reauthentication pause state, eliminating the need for manual CLI intervention or UPM scripts during a RADIUS outage.

Enhancement: Previously, the RADIUS reauthentication pause had to be triggered manually using a run-time CLI command. This enhancement allows the switch to automatically enter the pause state when a RADIUS timeout or failure is passively detected — without introducing active polling traffic to the network. While in the pause state, reauthentication timers for authenticated sessions are extended by a configurable maximum time, and new sessions can be assigned to a configured fallback VLAN. When RADIUS availability is passively restored through organic authentication traffic, the switch automatically resumes reauthentication timers.

Key Capabilities:

- Automatic pause state entry on passive detection of RADIUS server failure — no manual CLI command or UPM script required
- Configurable maximum pause duration to bound the outage window
- New sessions directed to a configured fallback VLAN during the pause period
- Automatic resumption of reauthentication timers when RADIUS availability is passively confirmed
- Manual override commands available to start or stop the pause state independently of automatic detection

New CLI Commands

Enable or disable automatic pause state entry on passive RADIUS failure detection:

```
configure netlogin radius reauthentication pause auto-start [enable | disable]
```

Manually start the RADIUS reauthentication pause:

```
start netlogin radius reauthentication pause
```

Manually stop the RADIUS reauthentication pause:

```
stop netlogin radius reauthentication pause
```

Modified CLI Command

The following command has been updated to support the new auto-start option:

```
configure netlogin radius reauthentication pause
```

Limitations

Because this feature uses passive detection only, detecting that the RADIUS server has come back online relies on organic authentication requests from clients. If no new clients attempt to authenticate, the switch remains in the paused state until the configured maximum timeout expires.

Supported Platforms

All Universal platforms.

Automatic Session Recovery After RADIUS Fallback VLAN Pause

Version 33.7.1 adds automatic re-authentication for sessions that were moved to the RADIUS fallback VLAN when the reauthentication pause was first triggered, restoring full network access without manual intervention after the RADIUS server becomes available again.

Enhancement: When the RADIUS reauthentication pause is activated, the initial session that detected the timeout is placed into the fallback VLAN to maintain basic connectivity. Previously, when the RADIUS server returned and the pause was lifted, that session remained in the fallback state until an administrator intervened manually. This enhancement introduces an automated re-trigger mechanism: the switch maintains a dynamic recovery list of MAC and 802.1X entries that were moved to the fallback VLAN as the pause trigger. When the pause ends — either by timer expiration or by a successful RADIUS response for another client — the switch automatically initiates re-authentication for all sessions on the recovery list.

Key Capabilities:

- Dynamic recovery list tracks sessions placed in the fallback VLAN at pause trigger time, across multiple ports simultaneously
- On pause removal, 802.1X sessions receive a port restart; MAC-Auth sessions receive a targeted re-authentication request
- Zero manual intervention required to restore full network access after a RADIUS outage

Modified CLI Command

The following command has been updated to support the automatic session recovery behavior:

```
configure netlogin radius reauthentication resume
```

Limitations

For 802.1X sessions, successful recovery depends on the client endpoint correctly responding to the port reset initiated by the switch.

Supported Platforms

All Universal platforms.

Configuration Recovery Removed from show inlets Output

Version 33.7.1 removes the configuration recovery endpoint references from the **`show inlets`** output.

Change: The NOS API configuration recovery endpoints (`/operation/system/config/recovery/:timeout` and `/operation/system/config/recovery/:cancel`), which previously triggered a reboot-based config rollback, have been removed. These endpoints are no longer listed in the **`show inlets`** output. Configuration rollback is now handled without a reboot as part of the Extreme Platform One configuration framework.

Supported Platforms

All Universal platforms.

EAPS FDB Flush Method Enhancement

The switch includes a new CLI command to optimize EAPS (Ethernet Automatic Protection Switching) convergence performance in large-scale deployments by providing configurable FDB flush methods.

Issue: In scaled EAPS environments with many shared-port links, protected VLANs, and EAPS domains, convergence events such as link failures or node reboots generate significant numbers of FDB flush operations. By default, the switch performs VLAN-and-port-based FDB flushing, which sends individual flush requests for each protected VLAN on the affected ring ports. As the number of protected VLANs and EAPS domains increases, the volume of FDB flush messages sent to hardware increases, potentially delaying convergence time.

Resolution: The switch supports port-based FDB flushing. This method performs a single port-level flush for all VLANs on the affected ring port, significantly reducing the number of hardware flush messages and improving convergence performance in large-scale EAPS deployments.

Key Capabilities:

- You can configure the FDB flush method to optimize convergence based on deployment scale: port-based for large-scale deployments and VLAN-and-port-based for standard deployments
- Port-based flushing performs a single flush operation for all VLANs on the affected ring port, reducing the total number of hardware flush messages
- VLAN-and-port-based flushing flushes FDB entries only for EAPS-protected VLANs on the affected ring ports and remains the default behavior
- Use port-based flushing for deployments with hundreds of protected VLANs and multiple EAPS shared-port links
- The configured flush method applies to all EAPS domains on the switch

New CLI Command

Configure the FDB flush method for EAPS convergence events:

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

Parameters

- `port-based` — Flushes FDB entries on EAPS ring ports for all VLANs. Use this method for deployments with hundreds of protected VLANs and multiple EAPS shared-port links to reduce the number of flush messages sent to hardware.
- `vlan-and-port-based` — Flushes FDB entries only for EAPS-protected VLANs on the affected ring ports. This is the default behavior and works well for deployments with moderate numbers of protected VLANs.

Example:

```
configure eaps flush-method port-based
```

Usage Guidelines

When you enable port-based flushing, the switch flushes FDB entries for all VLANs on the affected ring ports during convergence events. Temporary flooding occurs for both EAPS-protected and non-protected VLANs traversing those ring ports until the switch relearns forwarding entries. This trade-off is acceptable in large-scale deployments where convergence time is the priority.

VLAN-and-port-based flushing does not cause flooding for non-protected VLANs but generates more flush messages to hardware, making it less suitable for deployments with large numbers of protected VLANs and EAPS domains.

Supported Platforms

All platforms.

ExtremeCloud IQ Agent Logging Enhancements

The ExtremeCloud IQ Agent logging feature uses in-memory, buffer-based logging as the default mechanism instead of continuous on-disk file logging. This reduces storage consumption and improves operational efficiency.

Issue: ExtremeCloud IQ Agent records diagnostic logs to on-disk files with continuous logging during normal operation, even when no troubleshooting is in progress. This continuous disk logging consumes switch storage resources and makes real-time diagnostics difficult without manual file management.

Resolution: The logging feature makes buffer-based, in-memory logging the default. Recent diagnostic logs for both the Python ExtremeCloud IQ Agent application and the HiveAgent process are stored in EMS trace buffers that users can access directly from the CLI and view through `show tech-support` output. File-based, on-disk logging is disabled by default and available only as an on-demand debug-mode facility for extended troubleshooting sessions.

Key Capabilities:

- The ExtremeCloud IQ Agent process uses in-memory buffer-based logging with separate buffers for the Python application and the HiveAgent process
- File-based logging is suppressed by default, eliminating storage consumption for routine diagnostic operations
- Buffer logs are included in `show tech-support` output for support escalations and diagnostics
- File-based logging is available on demand at critical, error, warning, info, and debug severity levels using the debug CLI for targeted troubleshooting sessions
- Disabling file-based logging stops new writes but preserves previously captured on-disk log files
- The feature is backward compatible with older packages that do not support this logging control; the debug commands are accepted but do not change logging behavior

Behavior Change

After upgrading, ExtremeCloud IQ Agent file-based, on-disk logging is off by default, so the switch does not write new entries to on-disk log files. Any existing log files present on the switch at the time of upgrade are preserved and are not deleted or truncated. Recent activity remains available through the buffer-based trace feature. To resume or extend on-disk file-based logging during a troubleshooting session, enable file-based logging using the debug CLI.

Configuration and Usage

Buffer logs are viewed and controlled through EMS trace and debug CLI facilities. For detailed configuration procedures and complete CLI command syntax, refer to the user guide.

Supported Platforms

All platforms that support ExtremeCloud IQ Agent.

Fabric Attach Management VLAN Enhancements

The Fabric Attach feature includes conditional logic for automatic fallback management VLAN creation and server synchronization. This feature prevents disruption in environments where a different management VLAN is already in use while preserving zero-touch provisioning automation.

Issue: The automatic fallback management VLAN feature creates a VLAN in the FA proxy when the FA server does not provide one. In deployments where explicit FA management-VLAN configurations are already present or where zero-touch provisioning is not active, the automatic VLAN creation consumes unnecessary resources and can disrupt operations.

Resolution: The Fabric Attach feature applies conditional logic to fallback management VLAN creation and assignment. The fallback VLAN is created only when zero-touch provisioning is running and the FA server is not sending a management VLAN. Client port assignment and VLAN advertisement are conditional upon FA server approval of the Network Service Identifier (NSI) that the proxy requests. When the server rejects the NSI, the proxy retries while

zero-touch provisioning is active but does not apply configuration changes until the server approves the request.

Key Capabilities:

- Fallback management VLAN creation is exclusive to zero-touch provisioning scenarios: the switch creates the VLAN only when zero-touch provisioning is running and the FA server is not sending a management VLAN
- The switch assigns client ports to the management VLAN only when the FA server approves the Network Service Identifier (NSI) that the proxy requests
- When the FA server rejects the NSI, the proxy implements automatic retry attempts while zero-touch provisioning is active, but does not apply configuration changes during the retry period
- The switch does not advertise management VLAN information to FA clients when the server rejects the NSI, preventing unauthorized VLAN assignments
- Upon server approval of the NSI, the switch adds client ports to the management VLAN, advertises VLAN information to FA clients via LLDP, and persists the configuration until the next reboot or until the FA server provides an updated management VLAN assignment
- The `configure fabric attach management-vlan forward off` command allows operators to prevent VLAN plumbing on ports where FA clients are detected
- The feature is backward compatible with existing Fabric Attach configurations; no configuration changes are required for deployments that do not use zero-touch provisioning

Supported Platforms

All platforms.

FEC Statistics Detailed Display

Version 33.7.1 adds per-port Forward Error Correction (FEC) operational statistics for high-speed Ethernet interfaces, providing visibility into FEC health to help diagnose marginal links, optics issues, and cabling problems.

Enhancement: The OS previously displayed only FEC configuration state (enabled/disabled and clause) for high-speed ports such as 25G and 100G. Without operational counters, operators had no way to determine whether FEC was actively correcting errors or whether errors were exceeding correction capability. This enhancement exposes two per-port counters — corrected blocks and uncorrectable blocks — for both Clause 74 (BASE-R/Fire Code) and Clause 91/108 (Reed-Solomon) FEC modes. Counters are 64-bit monotonically increasing values that persist across link bounces and reset when FEC is disabled, when the FEC mode is toggled, or when explicitly cleared.

Key Capabilities:

- Corrected Blocks counter tracks FEC blocks or codewords that contained errors successfully corrected by the FEC decoder; a non-zero value indicates a marginal but functional link
- Uncorrectable Blocks counter tracks FEC blocks or codewords that exceeded correction capability and may result in traffic loss; any non-zero value indicates link degradation
- FEC counters can be cleared independently of other port counters using the new **`clear counters forward-error-correction`** command option, or cleared along with all counters via the existing **`clear counters`** command
- Ports that do not support FEC display blank counter values to preserve table readability

Modified CLI Command

The following command has been updated to include Corrected Blocks and Uncorrectable Blocks columns:

```
show ports {<port-list>} forward-error-correction {no-refresh}
```

New CLI Command

Clear FEC counters on all ports or a specified subset of ports:

```
clear counters forward-error-correction {ports {<port-list> | all}}
```

Limitations

Lane-level FEC statistics, SNMP/MIB access, and REST API/Telegraf support are not available in this release.

Supported Platforms

All platforms.

Flood Rate-Limit Burst Tolerance and Output Actions

The flood rate-limit feature includes configurable burst tolerance levels and optional out-of-profile actions. This allows operators to fine-tune packet buffer handling for broadcast, multicast, and unknown-destination MAC traffic on a per-port basis.

Issue: The flood rate-limit feature enforces a fixed buffer bucket size for controlling broadcast, multicast, and unknown-destination MAC traffic. Network administrators cannot customize burst behavior or define actions when traffic exceeds the rate limit, making it difficult to balance strict rate enforcement with brief traffic spikes in bursty network environments.

Resolution: The flood rate-limit feature includes two enhancements: (1) a configurable `burst-tolerance` parameter that controls buffer handling at different tolerance levels, and (2) optional `out-actions` that specify actions when traffic exceeds the configured rate limit, such as logging events, generating SNMP traps, or disabling the port.

Key Capabilities:

- Burst tolerance levels provide control over buffer size: default retains the existing behavior, low applies aggressive burst suppression, medium applies moderate burst tolerance, and high applies permissive burst tolerance
- The switch applies burst tolerance on a per-port basis to individual flood rate-limit traffic types: broadcast, multicast, and unknown-destmac
- Out-of-profile actions log rate-limit events, generate SNMP traps, or disable the port when traffic exceeds the configured rate limit
- Multiple out-of-profile actions can be combined in a single configuration
- Configuration persists across reboots and is included in configuration backups
- Display commands show the configured burst tolerance and current rate-limit status per port
- Existing flood rate-limit configurations remain functional with the default burst tolerance and no actions

Configuration Syntax

Configure flood rate-limit with burst tolerance and out-of-profile actions:

```
configure ports [<port_list> | <port_group>] rate-limit flood [broadcast | multicast |
unknown-destmac] <pps> { burst-tolerance [default | low | medium | high] } {out-actions
[{log} {trap} {disable-port}] }
```

Parameters

- `<pps>` — Rate limit in packets per second
- `burst-tolerance` — Buffer size control: default uses existing behavior, low applies strict suppression, medium applies moderate tolerance, high applies permissive tolerance
- `out-actions` — Actions when traffic exceeds the limit: log logs the event, trap sends an SNMP trap, disable-port disables the port

Example Configuration

```
configure ports 1 rate-limit flood broadcast 2000 burst-tolerance low out-actions log
```

This configuration limits broadcast traffic on port 1 to 2000 packets per second with low burst tolerance and logs any rate-limit violations.

Display Rate-Limit Configuration

View the configured burst tolerance and rate-limit status:

```
show port 1 information detail
```

Sample Output

```
Port: 1 Virtual-router: None Type: UTP ... Broadcast Rate: 2000 packets-per-second Burst
Tolerance: Low Multicast Rate: No-limit Unknown Dest Mac Rate: No-limit ...
```

Supported Platforms

All platforms.

IPP Enhancement: On-Demand Profile Rematch

Version 33.7.1 adds an on-demand rematch operation for Instant Port Profile (IPP), allowing matched devices to be re-evaluated against the current profile configuration.

Enhancement: Instant Port Profile matches connected devices to a profile device-type based on the best available information at the time of matching. Because matched devices remain associated with their original device-type even after switch state changes — such as a profile update adding new device-types, or a delayed LLDP neighborship that would have produced a better match than an initial MAC-learning match — the matched state can become inconsistent over time. The new **`run instant-port profile rematch`** command re-evaluates all currently matched devices and associates each with the most appropriate device-type based on current switch state.

New CLI Command

Trigger an on-demand rematch of all matched IPP devices:

```
run instant-port profile rematch
```

Supported Platforms

All platforms that support ExtremeCloud IQ.

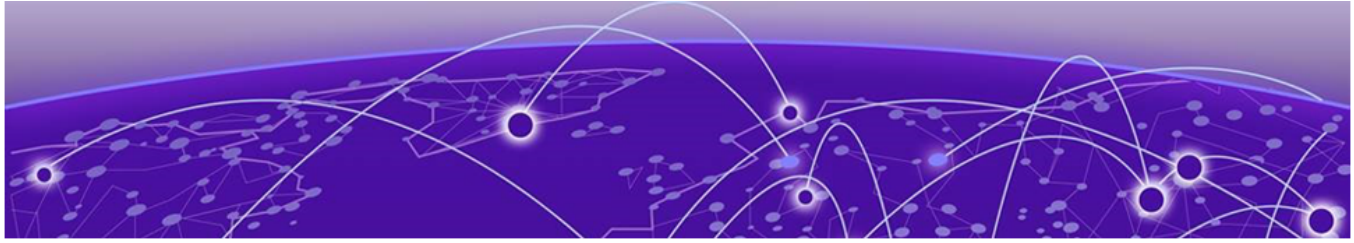

Changing the Network Operating System

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

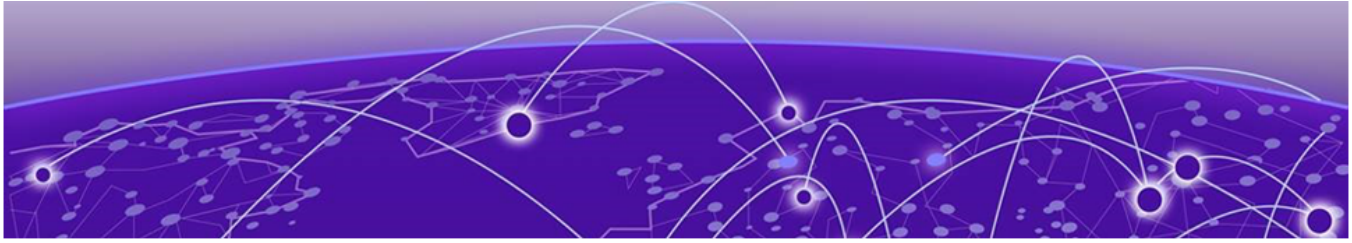

Extreme Hardware/Software Compatibility and Recommendation Matrices

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/.

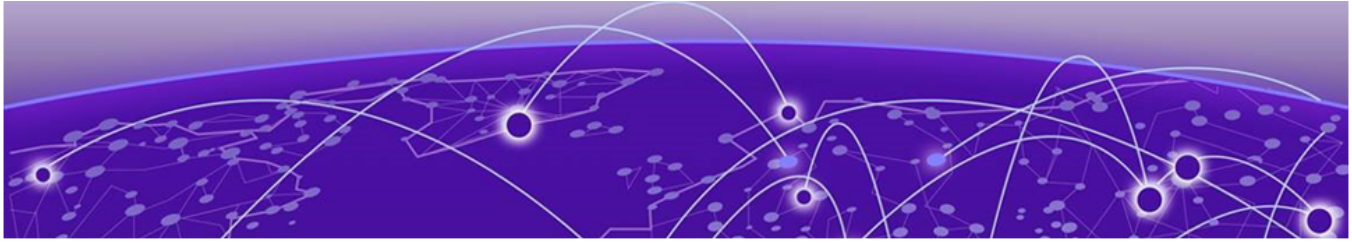

Compatibility with Extreme Management Center

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.

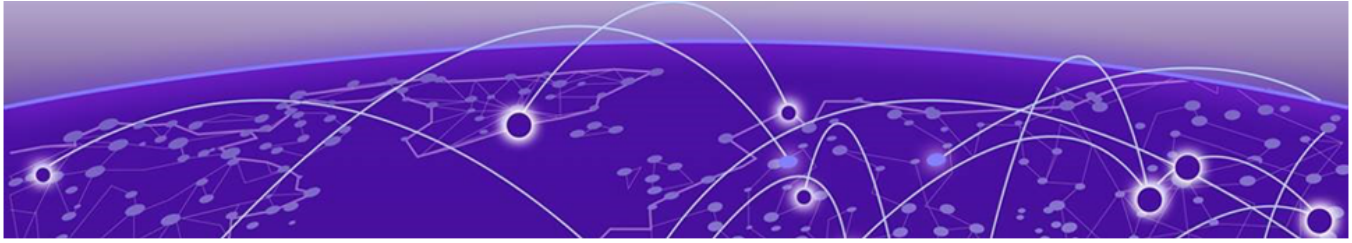

Supported MIBs

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.

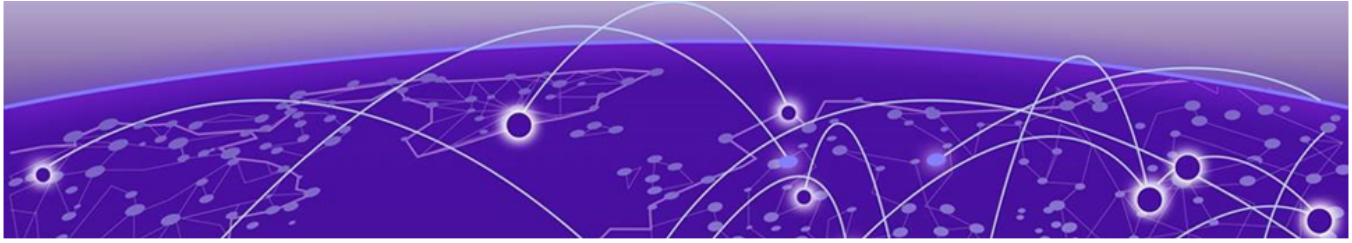

Tested Third-Party Products

The following third-party products have been tested.

Tested RADIUS Servers

The following RADIUS servers are fully tested:

- Microsoft—Internet Authentication Server
- Meetinghouse
- FreeRADIUS

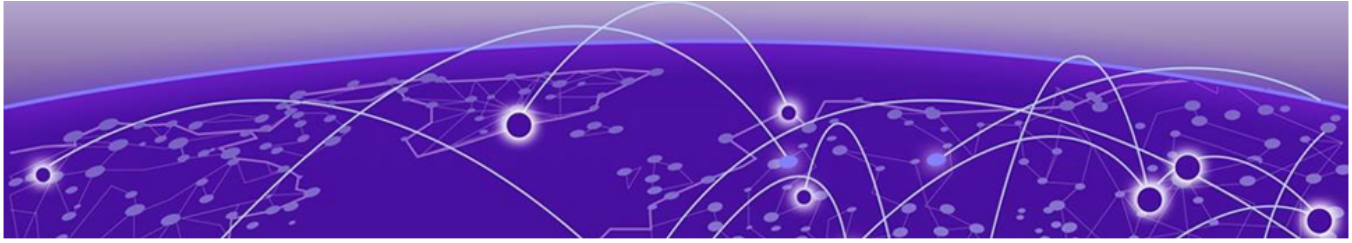

Extreme Switch Security Assessment

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

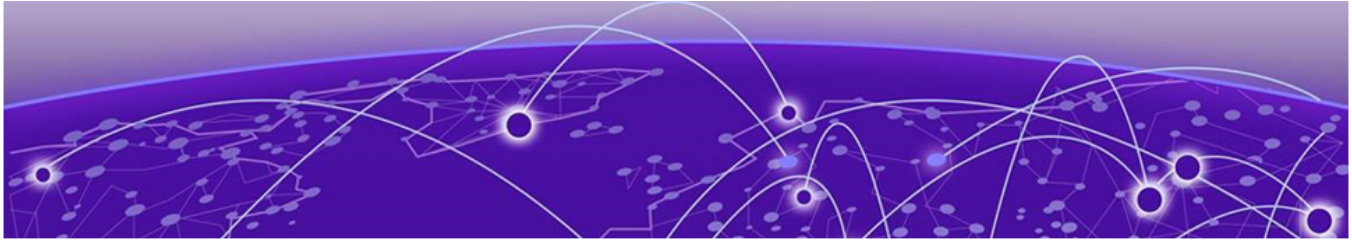

Limits

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 42
- Premier License Limits on page 78

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.

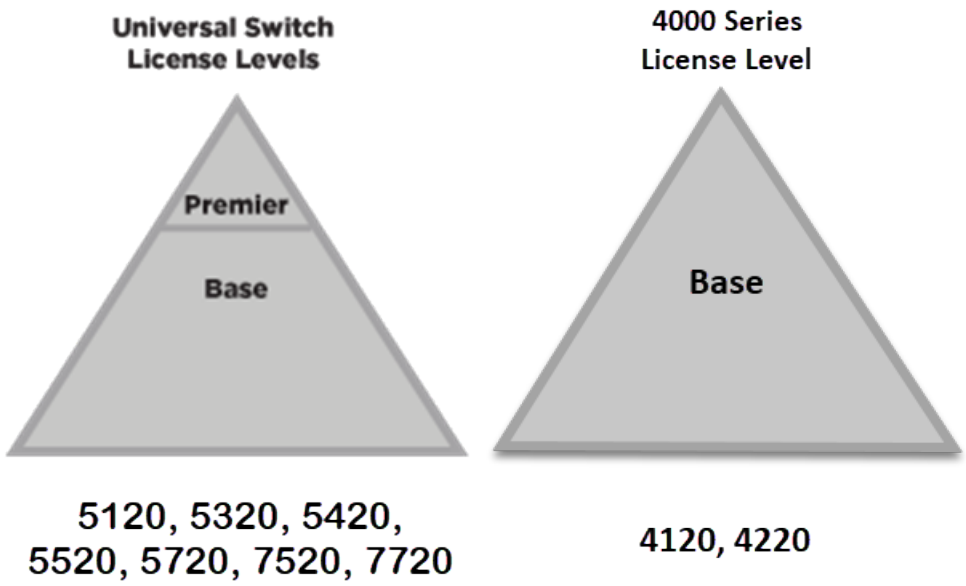

Figure 1: License Levels for Universal Switches

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

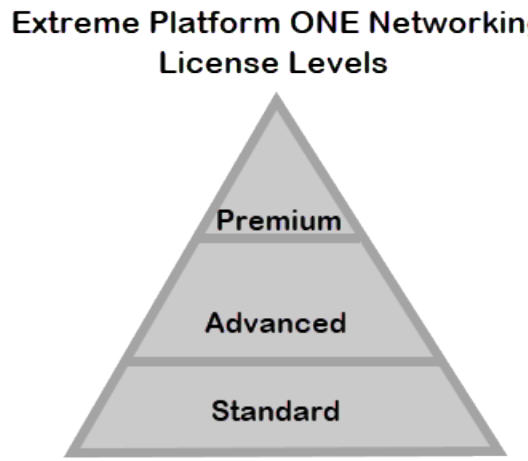

Figure 2: Extreme Platform ONE Networking License Levels

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 Networking license will perform the following actions:

- 5000 and 7000 series - activate Premier Universal license features

Extreme Platform ONE Networking 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.	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	5,000
	5720, 7520, 7720	10,000
	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, 5720-MXW, 7520, 7720	2,000
	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
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

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

Metric	Product	Limit
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, 5720-MXW, 7520, 7720	10,000
	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
OSPFv3 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes.	5520	3,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5720, 7520, 7720	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
PIM IPv4 Limits—maximum number of dynamic rendezvous points per multicast group.	All platforms	145
PIM IPv4 Limits—static rendezvous points.	All platforms	32

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

Metric	Product	Limit
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
Private VLANs—maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN.	4120, 4220, 5120, 5320, 5420, 5520, 5720 7520, 7720	36 71

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

Metric	Product	Limit
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 7520, 7720	960 1,024
Private VLANs—maximum number of private VLANs in an L2-only environment.	4120, 4220, 5120, 5320, 5420, 5520, 5720 7520, 7720	960 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 5320-16P, 5320-24T/P 5320-24T-4X-XT 4220, 5320-16P-2MXT-2X 4120, 5120	10,000 7,000 900 992 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 5120 5320-16P, 5320-24T/P 5320-16P-2MXT-2X 5320-24T-4X-XT	3,000 64 2,000 496 400
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 4120, 4220, 5120, 5320-24T/P, 5320-16P, 5320-24T-4X-XT	64 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 deafult)	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	5,000
	5720, 7520, 7720	10,000
	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, 5720-MXW, 7520, 7720	2,000
	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

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

Metric	Product	Limit
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
OSPFv3 external routes—recommended maximum number of external routes.	5520, 5720-MXW, 7520, 7720	10,000
	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	3,000
	5320 (except 5320-16P-2MXT-2X, 5320-24T-4X-XT), 5720, 7520, 7720	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
OSPFv3 virtual links—maximum number of OSPFv3 virtual links supported.	All platforms except 4120 and 4220	16

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

Metric	Product	Limit
PIM IPv4 (maximum interfaces)—maximum number of PIM active interfaces.	5320, 5420, 5520, 5720, 7520, 7720 5120	255 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 5120	3,000 (depends on policy file limits) 192
PIM IPv4 Limits—maximum number of multicast sources per group.	5320, 5420, 5520, 5720, 7520, 7720 5120	5,000 192
PIM IPv4 Limits—maximum number of dynamic rendezvous points per multicast group.	5320, 5420, 5520, 5720, 7520, 7720 5120	145 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 5120	255 30
PIM IPv6 limits—maximum number of multicast sources per group.	5320, 5420, 5520, 5720, 7520, 7720 5120	1,750 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 5120	3,000 (depends on policy file limits) 70
PIM IPv6 limits—maximum number of multicast groups per dynamic rendezvous points per multicast group.	5320, 5420, 5520, 5720, 7520, 7720 5120	64 20
PIM IPv6 limits—maximum number of secondary addresses per interface	5320, 5420, 5520, 5720, 7520, 7720 5120	70 30

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

Metric	Product	Limit
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	1 transparent clock
	7520-48Y, 7720-32C	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`.

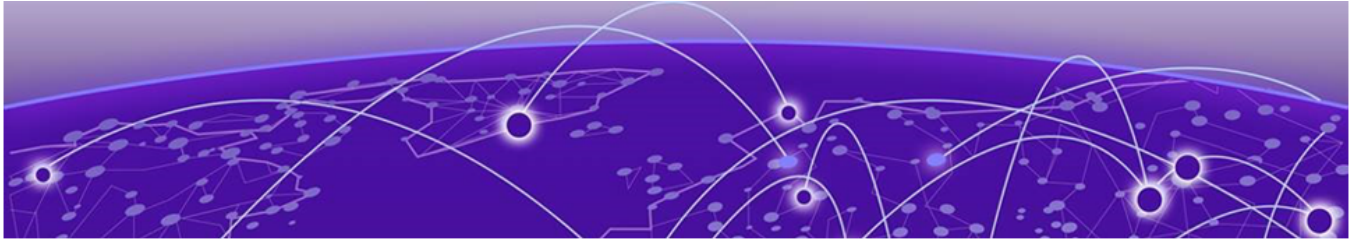

Open Issues, Known Behaviors, and Resolved Issues

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 33.7.1

Defect Number	Description
General	
CFD-16886	FDB entries are not learned from received IGMP reports when the switch is in software learning mode.

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 33.7.1

Defect Number	Description
EXOS-33947	LLDP packets sent by Windows 11 clients using destination MAC addresses `01:80:C2:00:00:00` and `01:80:C2:00:00:03` are not processed correctly as LLDP packets. There is no workaround at this time.
EXOS-34121	In a PIM Dense Mode multicast configuration, the pruned interface list in **`show pim ipv6 cache detail`** and **`show pim cache detail`** output may not be updated correctly on downstream nodes after an MSM failover. There is no workaround at this time.
EXOS-36691	PTPv2 Boundary Clock and Ordinary Clock functionality does not work correctly on 7x20 platforms when a loopback VLAN is configured as a PTP VLAN port. There is no workaround at this time.

Resolved Issues in Switch Engine Version 33.7.1-Patch1-26

The following issues were resolved in version 33.7.1.6-Patch-1-26. Version 33.7.1.6-Patch1-26 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, and 33.6.x.

Table 11: Resolved Issues and Feature Change Requests (CRs) in 33.7.1.6-Patch1-26

Key	Release Notes
General	
CFD-18051	Timezone configuration in "show configuration <devmgr>" output displays autodst even when it is not enabled explicitly.
CFD-18140	The command "show switch non-volatile-storage health-check" is restricted for non-admin users.
CFD-18193	The MRP process crashes when simultaneously enabling MVRP on an MLAG port on both MLAG peers.
CFD-18194	SFLOW sub-sampling is set to 2 even though all ports are configured with the same sample-rate.
CFD-18195	Password for user admin is not updated when attempting to change it via Web GUI or Chalet.
CFD-18230	The show switch management command displays image integrity checking as unknown instead of valid.
CFD-18231	Broadcast and Multicast counters are not incremented in "show port mgmt stat".

Table 11: Resolved Issues and Feature Change Requests (CRs) in 33.7.1.6-Patch1-26 (continued)

Key	Release Notes
CFD-18354	When the switch is in software-learning mode and with netlogin enabled, IGMP report packets with "Router-Alert" option are ignored.
EXOS-39960	In MLAG environment, packet blocking is not functioning correctly, resulting in duplicate packet forwarding and Layer 2 loops.
5320	
CFD-18240	Extended Diagnostics fails when a link is up at 100M.
VPEX	
CFD-18241	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.

Resolved Issues in Switch Engine Version 33.7.1-Patch1-11

The following issues were resolved in version 33.7.1.6-Patch-1-11. Version 33.7.1.6-Patch1-11 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, and 33.6.x.

Table 12: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in 33.7.1.6-Patch1-11

Defect Number	Description
General	
CFD-17028	PIM join messages were not processed by the RP when it receives (S, G, RPT) after the interface is pruned.
CFD-17056	NTP configuration was lost after reboot when DHCP is enabled on the VLAN.
CFD-17793	On switches without a dedicated OOB management port, a memory leak occurs in the nettools process whenever cloud-connector runs.
CFD-17796	VXLAN packets are processed by the CPU instead of being switched in hardware.
CFD-17823	CfgMgr process crashes during license check on Universal platforms if the switch has a user-created VR in the configuration.
CFD-17824	Improved log message when an Access profile denies access based on an implicit deny without an explicit deny entry in dynamic ACL.
CFD-17952	Added provision for port-based FDB-Flush in EAPS implementation, which can be used in scaled networks to minimize hardware programming load.
SummitStack	
CFD-17825	When one of the stack member is removed/rebooted, show fans display "operational at 0 RPM" instead of Empty.

Resolved Issues in Switch Engine

The following issues were resolved in version 33.7.1. Version 33.7.1 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, and 33.6.x.

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in 33.7.1

Defect Number	Description
General	
CFD-12612	An `IndexError: list assignment index out of range` exception occurs when executing commands that convert a VLAN tag to a VLAN name.
CFD-14876	Executing the **`show configuration`** command when the port display string contains German special characters causes a HAL process crash.
CFD-15767	An L3VPN route is not programmed correctly in hardware after certain network events.
CFD-15830	CLI sessions are not terminated based on the idle timeout configuration when the session was opened.
CFD-16145	The port bitmap of a port dynamically added to a VLAN is returned when the `dot1qVlanStaticTable` is polled.
CFD-16355	There is no CLI to configure the number of log entries stored on the switch.
CFD-16395	Pressing `q` in the **`show port utilization bandwidth`**, **`packets`**, or **`bytes`** command output does not quit and continues to print the next page.
CFD-16644	After adding a /32 route for an IP address that exists in the adjacency table, the switch does not accept new routes.
CFD-16710	The output of the **`show slpp guard ports`** command does not quit properly.
CFD-16972	A `can't open /proc/1705/cmdline: no such file` error message appears during switch bootup.
CFD-17056	The NTP configuration is lost after a reboot when DHCP is enabled on the VLAN.
CFD-17084	ExtremeCloud IQ device update fails on 5320 devices with the error `Collection for device is not complete.`
CFD-17087	An SNMP error occurs after upgrading to version 32.7.3.
CFD-17194	An error log message `Failed to set fec config on port 1, rv = -18` appears after rebooting the switch.
CFD-17256	BGP blackhole entry remains in the kernel after removing the corresponding NLRI entry from the BGP route policy.
EXOS-39140	IP multicast NAT does not work on a stack when the egress slot differs from the ingress slot.
EXOS-39261	Running cable diagnosis may results in unstable link and CRC errors.

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) in 33.7.1 (continued)

Defect Number	Description
EXOS-39401	ExtremeCloud IQ Site Engine does not save configuration changes made via the cloud connector after a CLI asterisk (unsaved changes indicator) is present.
EXOS-39547	DHCP relay cannot be configured from ExtremeCloud IQ.
EXOS-39621	When Instant-Port matches multiple device types where one has an Interlink action and another has an untagged VLAN action, the profile becomes invalid and fails to properly clean up actions when deleted.
Summit Stack	
CFD-16961	The **`show slot`** output does not display the currently selected partition on the master slot.
CFD-17226	The VM process is not initialized after a failover.
VPEX	
CFD-16670	SNMP general errors occur when enforcing VLAN configuration to a VPEX stack from ExtremeCloud IQ Site Engine.
4220 Series Switches	
CFD-17335	The **`show fan`** command returns the fan status as `Failed at 240 RPM`.
CFD-17395	The fan control algorithm assumes an incorrect fan model is installed on the 4220-24T-4X.
CFD-17662	Telemetry rules are not installed on re-configuration.
5520 Series Switches	
CFD-16239	When the 5520-48 is running version 33.4 or later, a connected dual-plane device sees 50% of all packets as CRC errors.
CFD-16913	OnePolicy does not filter incoming IGMP on the 5520.