



ExtremeXOS Release Notes

Software Version 16.1.3-Patch1-7

Copyright © 2016 All rights reserved.

Legal Notice

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

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

Trademarks

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

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

For additional information on Extreme Networks trademarks, please see:

www.extremenetworks.com/company/legal/trademarks/

Support

For product support, including documentation, visit: www.extremenetworks.com/documentation/

For information, contact:

Extreme Networks, Inc.

145 Rio Robles

San Jose, California 95134

USA

Table of Contents

| | |
|--|------------|
| Preface..... | 5 |
| Conventions..... | 5 |
| Providing Feedback to Us..... | 6 |
| Getting Help..... | 6 |
| Related Publications..... | 7 |
| Chapter 1: Overview..... | 9 |
| New and Corrected Features in ExtremeXOS 16.1.3-Patch1-7..... | 9 |
| New and Corrected Features in ExtremeXOS 16.1..... | 22 |
| ExtremeXOS Images for Summit X480 Series Switches..... | 41 |
| New Hardware Supported in 16.1.3..... | 42 |
| New Hardware Supported in ExtremeXOS 16.1..... | 42 |
| CLI Command Output Format of Ports Lists..... | 42 |
| Joint Interoperability Test Command (JITC) Compliance..... | 43 |
| Extreme Hardware/Software Compatibility and Recommendation Matrices..... | 43 |
| Compatibility with NetSight..... | 43 |
| Upgrading to ExtremeXOS..... | 43 |
| Downloading Supported MIBs..... | 43 |
| Tested Third-Party Products..... | 44 |
| Extreme Switch Security Assessment..... | 45 |
| Service Notifications..... | 45 |
| Chapter 2: Limits..... | 46 |
| Chapter 3: Open Issues, Known Behaviors, and Resolved Issues..... | 90 |
| Open Issues..... | 90 |
| Known Behaviors..... | 92 |
| Resolved Issues in ExtremeXOS 16.1.3-Patch1-7..... | 93 |
| Resolved Issues in ExtremeXOS 16.1.3-Patch1-4..... | 95 |
| Resolved Issues in ExtremeXOS 16.1.3-Patch1-2..... | 97 |
| Resolved Issues in ExtremeXOS 16.1.3..... | 100 |
| Resolved Issues in ExtremeXOS 16.1.2-Patch1-4..... | 101 |
| Resolved Issues in ExtremeXOS 16.1.2-Patch1-1..... | 103 |
| Resolved Issues in ExtremeXOS 16.1.2..... | 105 |
| Resolved Issues in ExtremeXOS 16.1..... | 111 |
| Chapter 4: ExtremeXOS Document Corrections..... | 126 |
| ACL Policy Redirect..... | 126 |
| ACL Ports Limits..... | 126 |
| Configure Sys-Recovery-level Slot Command Platform Availability..... | 127 |
| Configuring DHCP Binding..... | 127 |
| ELRP and QoS..... | 127 |
| ERPS/EAPS Failover Packet Loss in Stacking..... | 128 |
| Hardware Table Hash Algorithm..... | 128 |
| L2VPN Sharing Commands..... | 128 |
| Link Layer Discovery Protocol (LLDP)..... | 129 |
| Match Conditions Supported for OSPF Import Policy..... | 129 |
| MIB..... | 130 |
| MLAG PIM-SM *,G Forwarding Limitation..... | 130 |

| | |
|---|-----|
| NetLogin Limitation..... | 130 |
| NetLogin Local Authentication..... | 130 |
| PoE Power Delivery..... | 131 |
| PVLAN + EAPS and PVLAN + STP Recommendations..... | 131 |
| VRRP Guidelines..... | 132 |



Preface

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that are used throughout this guide.

Table 1: Notice Icons






| Icon | Notice Type | Alerts you to... |
|---|----------------|--|
|  | General Notice | Helpful tips and notices for using the product. |
|  | Note | Important features or instructions. |
|  | Caution | Risk of personal injury, system damage, or loss of data. |
|  | Warning | Risk of severe personal injury. |
|  | New | This command or section is new for this release. |

Table 2: Text Conventions

| Convention | Description |
|--|---|
| Screen displays | This typeface indicates command syntax, or represents information as it appears 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 with brackets, such as [Return] 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] |
| <i>Words in italicized type</i> | Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles. |

Platform-Dependent Conventions

Unless otherwise noted, all information applies to all platforms supported by ExtremeXOS software, which are the following:

- Summit® family switches
- SummitStack™
- ExtremeSwitching

When a feature or feature implementation applies to specific platforms, the specific platform is noted in the heading for the section describing that implementation in the ExtremeXOS command documentation. In many cases, although the command is available on all platforms, each platform uses specific keywords. These keywords specific to each platform are shown in the Syntax Description and discussed in the Usage Guidelines.

Terminology

When features, functionality, or operation is specific to a switch family, the family name is used. Explanations about features and operations that are the same across all product families simply refer to the product as the "switch."

Providing Feedback to Us

We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short [online feedback form](#). You can also email us directly at internalinfodev@extremenetworks.com.

Getting Help

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

- **Global Technical Assistance Center (GTAC) for Immediate Support**
 - **Phone:** 1-800-872-8440 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact
 - **Email:** support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- **GTAC Knowledge** — Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- **The Hub** — A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- **Support Portal** — Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) 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 Return Material Authorization (RMA) numbers

Related Publications

ExtremeXOS Publications

- *ACL Solutions Guide*
- *EMS Messages Catalog*
- *ExtremeXOS Command Reference Guide*
- *ExtremeXOS Feature License Requirements*
- *ExtremeXOS OpenFlow User Guide*
- *ExtremeXOS Quick Guide*
- *ExtremeXOS User Guide*
- *ExtremeXOS User Guide*
- *ExtremeXOS Legacy CLI Quick Reference Guide*
- *ExtremeXOS Release Notes*
- *Extreme Hardware/Software Compatibility and Recommendation Matrices*
- *Switch Configuration with Chalet*
- *Using AVB with Extreme Switches*

Hardware Documentation

- *E4G Series Routers Hardware Installation Guide*
- *Extreme Hardware/Software Compatibility and Recommendation Matrices*
- *Extreme Networks Pluggable Transceivers Installation Guide*
- *ExtremeSwitching X8 Series Switches Hardware Installation Guide*
- *Summit Family Switches Hardware Installation Guide for Switches Using ExtremeXOS 21.1*
- *Summit Family Switches Hardware Installation Guide for Switches Using ExtremeXOS 16 and Earlier*

Extreme Management Center™ Documentation

Extreme Management Center (EMC, formerly NetSight) documentation, including release notes, are available at: <https://extranet.extremenetworks.com/>. You must have a valid customer account to access this site.

Extreme Management Center online help is available from the **Help** menu in all EMC software applications. The online help provides detailed explanations of how to configure and manage your network using EMC software applications.

For complete regulatory compliance and safety information, refer to the document *Intel® Server Products Product Safety and Regulatory Compliance*.

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/software-licensing

1 Overview

New and Corrected Features in ExtremeXOS 16.1.3-Patch1-7
New and Corrected Features in ExtremeXOS 16.1
ExtremeXOS Images for Summit X480 Series Switches
New Hardware Supported in 16.1.3
New Hardware Supported in ExtremeXOS 16.1
CLI Command Output Format of Ports Lists
Joint Interoperability Test Command (JITC) Compliance
Extreme Hardware/Software Compatibility and Recommendation Matrices
Compatibility with NetSight
Upgrading to ExtremeXOS
Downloading Supported MIBs
Tested Third-Party Products
Extreme Switch Security Assessment
Service Notifications

These release notes document ExtremeXOS 16.1.3-Patch1-7, which adds a feature and resolves software deficiencies.

New and Corrected Features in ExtremeXOS 16.1.3-Patch1-7

This section lists the new and corrected features supported in the ExtremeXOS 16.1.3-Patch1-7 software:

ExtremeXOS Resiliency Enhancement for IPv4 and IPv6 Static Routes

This ExtremeXOS enhancement provides a resilient way to use Equal-Cost Multi-Path (ECMP) to load balance IPv4 or IPv6 traffic among multiple servers or other specialized devices. ExtremeXOS automatically manages the set of active devices using ECMP static routes configured with ping protection to monitor the health of these routes. Such servers or specialized devices do not require special software to support Bidirectional Forwarding Detection (BFD), or IP routing protocols such as OSPF, or proprietary protocols to provide keepalive messages. ExtremeXOS uses industry-standard and required protocols ICMP/ARP for IPv4, or ICMPv6/Neighbor Discovery for IPv6, to accomplish the following automatically:

- Initially verify devices and activate their static routes, without waiting for inbound user traffic, and without requiring configuration of device MAC addresses.
- Detect silent device outages and inactivate corresponding static routes.
- Reactivate static routes after device recovery, or hardware replacement with a new MAC address.

ExtremeXOS currently supports similar protection and resiliency using Bidirectional Forwarding Detection (BFD) on IPv4 and IPv6 static routes. However, BFD can only be used when the local and remote device both support BFD.

New Commands ExtremeXOS Resiliency Enhancement for IPv4 and IPv6 Static Routes

The following are new commands for the ExtremeXOS Resiliency Enhancement for IPv4 and IPv6 Static Routes feature.



configure iproute add protection

```
configure iproute add [default | ipv4_or_ipv6_network] gateway { protection [ bfd | ping | none ] }
```

Description

Configures protection and resiliency on IPv4 and IPv6 static routes.

Syntax Description

| | |
|-----------------------------|---|
| default | Default route. |
| <i>ipv4_or_ipv6_network</i> | IPv4 or IPv6 network address. |
| <i>gateway</i> | Gateway IP address. |
| protection | Selects the type of protection on this route (default is none). |
| bfd | Enables BFD protection on this route. |
| ping | Enables ICMP ping protection on this route. |
| none | Disables all protection on this route (default). |

Default

No protection is the default.

Usage Guidelines

For static routes configured with protection type ping, static routes are initially down. Static routes become "up" for each configured gateway/device IP when a timely ICMP Echo Reply is received from that IP within the configured ping interval. Static routes transition from up to down when no timely reply is received for the configured number of missed intervals. Severely delayed ICMP Echo Replies are ignored if received after the configured interval time elapses, because a new ICMP Echo Request has already been sent. Static routes with ping protection need not be ECMP routes. Thus when a device is unresponsive, a different route with a higher cost or shorter prefix length can route packets elsewhere.

For more information about static routes configured with protection type "BFD," see documentation regarding the `configure iproute add bfd` command.

The protection type (BFD, ping or none) for an existing static route can be changed dynamically without deleting the route. To change the protection type, simply re-add an existing static route with a different protection type.

Example

The following example adds a static route for 100.0.0.0/24 with ping health check monitoring to gateway IP 1.2.3.4.

```
configure iproute add 100.0.0.0/24 1.2.3.4 protection ping
```

ExtremeXOS initiates ping health check monitoring to the adjacent device with IP address 1.2.3.4. The route for 100.0.0.0/24 is protected, meaning if ping responses are received from 1.2.3.4 in a timely manner, the static route for 100.0.0.0/24 to 1.2.3.4 is “up” in the routing table. If no ping response is received in a timely manner, the route is down.

In an example with ECMP, assuming `enable iproute sharing`:

```
configure iproute add 100.0.0.0/24 1.2.3.5 protection ping
```

If ping responses are received by both 1.2.3.4 and 1.2.3.5, IP packets destined to subnet 100.0.0.0/24 are Layer-3 load balanced by hardware between 1.2.3.4 and 1.2.3.5. If for example, no ping response is received from 1.2.3.4 in a timely manner, IP packets destined to 100.0.0.0/24 are sent only to 1.2.3.5. Later, upon receiving a ping response from 1.2.3.4, packets are load balanced again.

History

This command was first available in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

All platforms with an Edge license or greater.



configure iproute protection ping interval

```
configure iproute {ipv4 | ipv6} protection ping interval seconds miss misses
```

Description

Configures the desired interval between pings, and number of misses for ping protection of IPv4 and IPv6 static routes.

Syntax Description

| | |
|-------------------|---|
| ipv4 | Designates IPv4 settings (default). |
| ipv6 | Designates IPv6 settings. |
| protection | Configures route protection settings. |
| ping | Configures static route ping protection interval and number of misses. |
| interval | Number of seconds between pings to protected gateways. Ping response must be received within configured interval. |
| <i>seconds</i> | Number of seconds between pings to protected gateways. Range is 1-600. Default is 2. |

| | |
|---------------|--|
| miss | Number of pings with no response before associated routes are considered down. |
| <i>misses</i> | Number of pings with no response before associated routes are considered down. Range is 2–255. Default is 3. |

Default

If not specified, IPv4 is the default, and:

- Interval = 2 seconds.
- Misses = 3.

Usage Guidelines

At the configurable interval, each unique gateway or device IP address configured for static route ping protection is sent an ICMP or ICMPv6 Echo Request if the ARP or Neighbor cache entry already has the IP->MAC binding. An ARP or Neighbor Solicitation is sent if the IP->MAC binding is unknown, and upon receiving a response, the ICMP Echo Request is sent.

The desired interval between pings, and number of misses, can be configured independently for IPv4 and IPv6.

Example

The following example sets for IPv4 a ping interval of 3 seconds and number of missed pings to 5:

```
configure iproute ipv4 protection ping interval 3 miss 5
```

History

This command was first available in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

All platforms with an Edge license or greater.



enable iproute protection ping

```
enable iproute {ipv4 | ipv6} protection ping
```

Description

Globally enables ping protection for static routes added with ping protection for IPv4 and IPv6.

Syntax Description

| | |
|-------------------|---|
| ipv4 | Specifies IPv4 (default). |
| ipv6 | Specifies IPv6. |
| protection | Enables route protection. |
| ping | Globally enables ping protection for static routes added with ping protection (default is enabled). |

Default

Enabled is the default. If not specified, IPv4 is the default.

Example

The following example enables ping protection for static routes added with ping protection for IPv4:

```
enable iproute ipv4 protection ping
```

History

This command was first available in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

All platforms with an Edge license or greater.



disable iproute protection ping

```
disable iproute {ipv4 | ipv6} protection ping
```

Description

Globally disables ping protection for static routes added with ping protection for IPv4 and IPv6. Routes are up in the routing table, and ping health check monitoring is not performed..

Syntax Description

| | |
|-------------------|--|
| ipv4 | Specifies IPv4 (default). |
| ipv6 | Specifies IPv6. |
| protection | Disables route protection. |
| ping | Globally disables ping protection for static routes added with ping protection (default is enabled). |

Default

Enabled is the default. If not specified, IPv4 is the default.

Example

The following example disables ping protection for static routes added with ping protection for IPv4:

```
disable iproute ipv4 protection ping
```

History

This command was first available in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

All platforms with an Edge license or greater.



show iproute protection ping

```
show iproute {ipv4 | ipv6} protection ping {v4_or_v6_gateway} {vr vr_name}
{detail}
```

Description

Displays the current ping protection state for IPv4 or IPv6 routes, along with the number of protected routes. The information can be displayed per gateway, or for all gateways in a given VR.

Syntax Description

| | |
|-------------------------|---|
| ipv4 | Specifies IPv4 routes (default). |
| ipv6 | Specifies IPv6 routes. |
| protection ping | Specifies protection ping information. |
| v4_or_v6_gateway | Name of the specific IPv4 or IPv6 gateway to show protected routes for. |
| vr vr_name | Shows all protected routes for the specified VR. |
| detail | Displays timestamps of the last time each gateway transitioned to the "up" state and to the "down" state. |

Default

If not specified, IPv4 is the default.

Example

The following example shows all ping protected IPv4 routes only:

```
# show iproute protection ping
Ping protection for static IPv4 routes : Enabled
Interval : 2
Misses : 3
```

```
Gateway      State  # Routes  VR
-----
1.2.3.4      Down   1         VR-Default
1.2.3.5      Up     1         VR-Default
```

The following example shows ping protected IPv6 routes only:

```
# show iproute ipv6 protection ping
Ping protection for static IPv6 routes : Enabled
Interval : 2
Misses : 3
```

```
Gateway      State  # Routes  VR
-----
2000:2000:2000:2000:2000:2000:2000:2001  Up     2         VR-Default
3000::2      Down   2         VR-Default
```

The following example shows ping protected IPv6 routes with gateway up/down transition timestamp information:

```
# show iproute ipv6 protection ping detail
Ping protection for static IPv6 routes : Enabled
      Interval : 2
      Misses   : 3

Gateway: 2000:2000:2000:2000:2000:2000:2000:2001
State:   Up
# Routes: 2
VR:     VR-Default
Last Up: Thu Mar 31 22:44:12 2016
Last Down: Thu Mar 31 22:42:22 2016

Gateway: 3000::2
State:   Down
# Routes: 2
VR:     VR-Default
Last Up: --
Last Down: Thu Mar 31 22:42:22 2016
```

History

This command was first available in ExtremeXOS release 16.1.3-Patch1-7.

Platform Availability

All platforms with an Edge license or greater.

Revised Commands ExtremeXOS Resiliency Enhancement for IPv4 and IPv6 Static Routes

The following are revised commands for the ExtremeXOS Resiliency Enhancement for IPv4 and IPv6 Static Routes feature.

- The configuration settings for static route ping protection enable/disable, interval, and misses for IPv4 appear in the [show ipconfig](#) on page 15 and for IPv6 in the [show ipconfig ipv6](#) on page 17 commands, respectively.
- A new route flag letter "I" appears in the [show iproute](#) on page 18 and [show ipconfig ipv6](#) on page 17 commands to indicate static routes with ICMP ping protection. Flag letter "I" uses the same column as flag letter "b" because BFD and ping protection are mutually exclusive. If the route flags also show "U" for Up, then ping protection detected the gateway is up.

show ipconfig

```
show ipconfig {ipv4} {vlan vlan_name}
```

Description

Displays configuration information for one or more VLANs in the current VR context.

Syntax Description

| | |
|------------------------|------------------------|
| <code>vlan_name</code> | Specifies a VLAN name. |
|------------------------|------------------------|

Default

N/A.

Usage Guidelines

If no VLAN information is specified, then global IP configuration is displayed. Otherwise, specific VLAN information is displayed.

Example

```
# show ipconfig
  Use Redirects : Disabled
  IpOption LSRR : Enabled
  IpOption SSRR : Enabled
  IpOption RR   : Enabled
  IpOption TS   : Enabled
  IpOption RA   : Enabled
  Route Sharing : Disabled
  Route Compression : Enabled
  Originated Packets : Don't require ipforwarding
  IP Fwding into LSP : Disabled
  Max Shared Gateways : Current: 4   Configured: 4
  Route Sharing Hash : CRC Lower
  Ping Protection : Enabled for static routes   Interval: 2   Misses: 3

  IRDP:
    Advertisement Address: 255.255.255.255   Maximum Interval: 600
    Minimum Interval: 450   Lifetime: 1800   Preference: 0

Interface      IP Address      Flags      nSIA

Flags: (A) Address Mask Reply Enabled (B) BOOTP Enabled
      (b) Broadcast Forwarding Enabled
      (D) Duplicate address detected on VLAN, (E) Interface Enabled
      (f) Forwarding Enabled (g) Ignore IP Broadcast Enabled
      (h) Directed Broadcast Forwarding by Hardware Enabled
      (I) IRDP Advertisement Enabled, (M) Send Parameter Problem Enabled
      (m) Multicast forwarding Enabled, (n) Multinetted VLAN
      (nSIA ) Number of Secondary IP Addresses
      (P) Send Port Unreachables Enabled, (R) Send Redirects Enabled
      (r) Unicast Reverse Path Enabled on at least one port of the VLAN
      (t) Tentative address, (T) Time Stamp Reply Enabled
      (u) Send Unreachables Enabled, (U) Interface Up
      (v) VRRP Enabled, (X) Send Time Exceeded Enabled
```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

This command changed to display information for the current VR context in ExtremeXOS 12.5.

Configuration settings for static route ping protection enable/disable, interval, and misses added in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

This command is available on all platforms.

show ipconfig ipv6

show ipconfig ipv6 {**vlan** *vlan_name* | **tunnel** *tunnel_name*}

Description

Displays configuration information for one or more interfaces in the current VR context.

Syntax Description

| | |
|--------------------|------------------------------------|
| <i>vlan_name</i> | Specifies an IPv6 configured VLAN. |
| <i>tunnel_name</i> | Specifies an IPv6 tunnel. |

Default

N/A.

Usage Guidelines

If no interface is specified, then global IPv6 configuration is displayed. Otherwise, specific interface(s) will be displayed. Global IPv6 configuration information includes:

- IPv6 address/netmask/etc.
- IPv6 forwarding information/IPv6 multicast forwarding information

Example

The following example displays configuration information on a VLAN "accounting":

```
# show ipconfig ipv6 accounting
Router Interface on accounting is enabled.  MTU: 1500
  IPv6 Hop Limit: 64
  Locally registered unicast addresses:
    2002::1/64 (Tentative)
    fe80::204:96ff:fe1e:ec00%accounting/64 (Tentative)
  Flags:
    IPv6 Forwarding: NO   Accept recvd RA: NO
    Send redirects: YES  Accept redirects: NO
```

The configuration settings for static route ping protection enable/disable, interval, and misses also appear:

```
# show ipconfig ipv6
Route Sharing           : Enabled
Route Compression      : Enabled
ICMP Redirect for Fast Path : Disabled
Max Shared Gateways    : Current: 16  Configured: 16
Route Sharing Hash     : CRC Lower
Ping Protection        : Enabled for static routes  Interval: 2
Misses: 3
```

History

This command was first available in ExtremeXOS 11.2.

This command changed to display information for the current VR context in ExtremeXOS 12.5.

The show output for **max-gateways** was added in ExtremeXOS 15.3.

Configuration settings for static route ping protection enable/disable, interval, and misses added in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature.

show iproute

```
show iproute {ipv4} {priority | vlan vlan_name | permanent | ip_address netmask |
summary} {multicast | unicast} {vr vrname}}
```

Description

Displays the contents of the IP routing table or the route origin priority.

Syntax Description

| | |
|-------------------|--|
| priority | Displays the priority values for each route origin type. |
| <i>vlan_name</i> | Specifies a VLAN name. |
| permanent | Specifies permanent routing. |
| <i>ip_address</i> | Specifies an IP address. |
| <i>netmask</i> | Specifies a subnet mask. |
| summary | Displays summary information. |
| multicast | Displays information for IPv4 multicast routes only. |
| unicast | Displays information for IPv4 unicast routes only. |
| <i>vrname</i> | Specifies a VR or VRF. |

Default

N/A.

Usage Guidelines

A c flag in the Flags column indicates a compressed route resulting from enabling compression using the enable iproute compression command. The total number of compressed routes is also shown.

All routes that are provided to the FIB display the f flag.

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example displays detailed information about all IP routing:

```
Switch.3 # show iproute
Ori Destination Gateway Mtr Flags VLAN Duration
d 2.2.0.0/16 2.2.2.3 1 -----um--- v2 2d:10h:17m:41s
#d 3.2.2.0/24 3.2.2.23 1 U-----um--f jim_igmp 1d:19h:49m:49s
```

```

d 3.3.3.0/24      3.3.3.1      1  -----um--- v3      2d:10h:17m:49s
d 4.4.4.0/24      4.4.4.2      1  -----um--- v4      2d:10h:17m:5s
*hm 192.168.23.1/32 192.168.23.1 1  UGHD---u---f- vlan23 0d:0h:16m:5s
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (il) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPEExt, (mbi) MBGPInter, (mp) MPLS Lsp
(mo) MOSPF (o) OSPF, (o1) OSPFExt1, (o2) OSPFExt2
(oa) OSPFIntra, (oe) OSPFAsExt, (or) OSPFInter, (pd) PIM-DM, (ps) PIM-SM
(r) RIP, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (hm) Host-mobility, (un)
UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (b) BFD protection requested, (B) BlackHole, (c) Compressed, (D)
Dynamic,
      (f) Provided to FIB, (G) Gateway, (H) Host Route,
      (I) ICMP ping protection requested, (l) Calculated LDP LSP,
      (L) Matching LDP LSP, (m) Multicast, (p) BFD protection active,
      (P) LPM-routing, (R) Modified, (s) Static LSP, (S) Static,
      (t) Calculated RSVP-TE LSP, (T) Matching RSVP-TE LSP, (u) Unicast,
      (U) Up, (3) L3VPN Route.

```

```

Mask distribution:
1 routes at length 16      3 routes at length 24
Route Origin distribution:
4 routes from Direct
Total number of routes = 4
Total number of compressed routes = 0

```

History

This command was first available in ExtremeXOS 10.1.

The **ipv4** keyword was added in ExtremeXOS 11.2.

The **c** flag was added in ExtremeXOS 12.0.

The **f** flag was added in ExtremeXOS 12.2.2.

The **hm** flag was added in ExtremeXOS 16.2

Support added for host mobility in ExtremeXOS 16.2.

"**l**" flag added showing ICMP ping protection in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

This command is available on all platforms.

show iproute ipv6

```

show iproute ipv6 {priority | vlan vlan_name | tunnel tunnel_name | ipv6Netmask |
summary {multicast | unicast}} {vr vr_name}}

```

Description

Displays the contents of the IPv6 routing table.

Syntax Description

| | |
|--------------------|--|
| priority | Displays the priority values for each route origin type. |
| <i>vlan_name</i> | Specifies a VLAN name. |
| <i>tunnel_name</i> | Specifies a tunnel name. |
| <i>ipv6Netmask</i> | Specifies an IPv6 address/prefix length. |
| summary | Specifies summary information. |
| <i>vr_name</i> | Specifies a VR or VRF. |

Default

N/A.

Usage Guidelines

If you do not specify a VR or VRF, the command applies to the current VR context.

Example

The following example displays detailed information about all IPv6 routing:

```
Switch.18 # show iproute ipv6
Ori Destination                               Mtr  Flags          Duration
Gateway                                       Interface
*hm 2000::/128                               1    UGHD---u---f- 0d:0h:0m:
7s
#d 2001:db8::/64                              1    U-----um--f 0d:0h:5m:
31s
2001:db8::52                                  ixia
#or 2001:db8:2:78::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:79::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:7a::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:7b::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:7c::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:7d::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:7e::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:7f::/64                        50   UG-D---um--f 0d:0h:0m:
1s
fe80::200:40ff:feba:a38e                      ixia
#or 2001:db8:2:80::/64                        50   UG-D---um--f 0d:0h:0m:
1s
```

```

fe80::200:40ff:feba:a38e          ixia
#or 2001:db8:2:81::/64           50   UG-D---um--f 0d:0h:0m:
ls
fe80::200:40ff:feba:a38e          ixia
#d fe80::%ixia/64                1    U-----um--f 0d:0h:5m:
31s
fe80::204:96ff:fe27:8697          ixia
Origin(Ori): (b) BlackHole, (be) EBGp, (bg) BGP, (bi) IBGP, (bo) BOOTP
(ct) CBT, (d) Direct, (df) DownIF, (dv) DVMRP, (e1) ISISL1Ext
(e2) ISISL2Ext, (h) Hardcoded, (i) ICMP, (i1) ISISL1 (i2) ISISL2
(is) ISIS, (mb) MBGP, (mbe) MBGPExt, (mbi) MBGPInter, (ma) MPLSIntra
(mr) MPLSInter, (mo) MOSPF (o) OSPFv3, (o1) OSPFv3Ext1, (o2) OSPFv3Ext2
(oa) OSPFv3Intra, (oe) OSPFv3AsExt, (or) OSPFv3Inter, (pd) PIM-DM, (ps) PIM-SM
(r) RIPng, (ra) RtAdvrt, (s) Static, (sv) SLB_VIP, (hm) Host-mobility, (un)
UnKnown
(*) Preferred unicast route (@) Preferred multicast route
(#) Preferred unicast and multicast route
Flags: (b) BFD protection requested, (B) BlackHole, (c) Compressed, (D)
Dynamic,
      (f) Provided to FIB, (G) Gateway, (H) Host Route,
      (I) ICMP ping protection requested, (l) Calculated LDP LSP,
      (L) Matching LDP LSP, (m) Multicast, (p) BFD protection active,
      (P) LPM-routing, (R) Modified, (s) Static LSP, (S) Static,
      (t) Calculated RSVP-TE LSP, (T) Matching RSVP-TE LSP, (u) Unicast,
      (U) Up, (3) L3VPN Route.
Mask distribution:
12 routes at length 64
Route Origin distribution:
2 routes from Direct              10 routes from OSPFv3Inter
Total number of routes = 12
Total number of compressed routes = 0

```

The following example displays the IPv6 route origin priority:

```

# show iproute ipv6 priority
Direct 10
Blackhole 50
Static 1100
ICMP 1200
OSPFv3Intra 2200
OSPFv3Inter 2300
Isis 2350
IsisL1 2360
IsisL2 2370
RIPng 2400
OSPFv3AsExt 3100
OSPFv3Ext1 3200
OSPFv3Ext2 3300
IsisL1Ext 3400
IsisL2Ext 3500

```

History

This command was first available in ExtremeXOS 11.2.

The hm flag was added in ExtremeXOS 16.2

Support added for host mobility in ExtremeXOS 16.2.

"I" flag added showing ICMP ping protection in ExtremeXOS 16.1.3-Patch1-7.

Platform Availability

This command is available on the platforms listed for the IPv6 unicast routing feature

New and Corrected Features in ExtremeXOS 16.1

This section lists the new and corrected features supported in the ExtremeXOS 16.1 software:

RADIUS Authentication and Authorization Enhancements

The RADIUS client software sends authentication requests using standard mechanisms for PAP, CHAP (RFC 2865 (13)) and EAP (RFC 3579 (12)).

This feature introduces authentication retransmission algorithm capability, which uses two retransmission algorithms in combination: Back-off Round Robin, and simple Round Robin. These retransmission algorithms provide server redundancy.

Eight authentication servers are now supported.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches
- E4G-200 and E4G-400 cell site routers

New CLI Commands

```
configure radius algorithm [standard | round-robin]
```

```
configure radius algorithm [standard | round-robin]
```

```
unconfigure radius-accounting [server index]
```

```
unconfigure radius [server index]
```

Changed CLI Commands

Changes are underlined.

```
configure radius {mgmt-access | netlogin} [primary | secondary | index] server  
[host_ipaddr | host_ipV6addr | hostname] {udp_port} client-ip [client_ipaddr |  
client_ipV6addr] {vr vr_name} {shared-secret {encrypted} secret}
```

```
configure radius [primary | secondary index] shared-secret  
{encryptedencrypted_secret | secret}
```

```
configure radius-accounting {mgmt-access | netlogin} [primary | secondary | index]  
server [host_ipaddr | host_ipV6addr | hostname] {udp_port} client-ip  
[client_ipaddr | client_ipV6addr] {vr vr_name} {shared-secret {encrypted} secret}
```

```
show radius {mgmt-access | netlogin} {primary | secondary | index}
```

```
show radius-accounting {mgmt-access | netlogin} {primary | secondary | index}
```

ONEPolicy

ONEPolicy allows you create profiles for securing and provisioning network resources based upon the role the user or device plays within the enterprise. By first defining the user or device role, network resources can be tailored to a specific user, system, service, or port-based context by configuring and assigning rules to the policy role. A policy role can be configured for any combination of Class of Service, VLAN assignment, classification rule precedence, or default behavior based upon L2, L3, and L4 packet fields. Hybrid authentication allows either policy or dynamic VLAN assignment, or both, to be applied through RADIUS authorization.

Supported Platforms

- Summit X450-G2
- Summit X460-G2
- Summit X670-G2
- Summit X770

Limitations

- ExtremeXOS only allows policy to be enabled if all the devices in the stack support policy. At the time of configuration a switch provisions the lowest common denominator of functionality. If a switch attempts to join the stack after policy is enabled, it must be able to support the existing level of functionality or it is not allowed to participate in policy.
- Only 'macdest', 'macsource', or 'port' policy rules can be applied to QinQ (that is, double-tagged) packets received on an untagged VMAN port.

New Commands

Policy Commands

```
enable policy
```

```
disable policy
```

```
configure policy vlanauthorization port [port_list | all] [{enable | disable}
{tagged | untagged}]
```

```
unconfigure policy vlanauthorization [{port port_list} | all]
```

```
configure policy invalid action [default-policy | drop | forward]
```

```
configure policy port ports admin-id admin_id
```

```
configure policy profile profile_index {name name } {pvidpvid } {pvid-status
pvid_status } {cos cos} {cos-status cos_status } {egress-vlans egress_vlan_list }
{untagged-vlans untagged_vlans } {append | clear} {tcioverwrite tci_overwrite }
```

```

configure policy rule admin-profile macsource macsource | portport ] {maskmask }
{ port-string [port_string | all]} {storage-type [non-volatile | volatile]}
{adminpid admin_pid } {tci-overwrite}

configure policy rule profile_index | etherether | ip6destip6dest | ipdestipdest
| ipfrag | ipprotoipproto | ipsourcesocketipsourcesocket | iptosiptos |
ipttlipttl | macdestmacdest | macsource macsource | portport |
tcpdestportIPtcpdestportIP | tcpsourceportIPtcpsourceportIP | udpdestportIP
udpdestportIP | udpsourceportIP udpsourceportIP] {maskmask } {port-string
[port_string | all]} {storage-type [non-volatile | volatile]} {drop | forward}
{coscos}

configure policy mactable [response [tunnel | policy | both] | vlan_list
profile_index]

show policy state

show policy allowed-type ports {detail}

show policy capability

show policy dynamic override

show policy invalid [all | {action} {count}]

show policy mactable

show policy profile

show policy rule {all | {profile-index profile_index | admin-profile} ether
{etner} ipdest {ipdest} | ipfrag | ipproto {ipproto} | ipsource {ipsource} |
iptos {iptos} | ipttl {ipttl} macdest {macdest } | macsource {macsource} | port
{port} | tcpdestportIP {tcpdestportIP} | tcpsourceportIP {tcpsourceportIP} |
udpdestportIP {udpdestportIP} | udpsourceportIP {udpsourceportIP}} {maskmask}
{port-string [port_string | all]} {storage-type [non-volatile | volatile]} {drop
| forward} {coscos | admin-pidadmin_pid }}{detail | wide}

show policy vlanauthorization {port port_list}

unconfigure policy all-rules

unconfigure policy mactable [response | vlan_list]

unconfigure policy profile [all | profile_index]

unconfigure policy vlanauthorization [enable | disable]

unconfigure policy rule [profile_index] [all-pid-entries | ipfrag | icmp6type |
icmptype | ip6dest | ipdest | ipproto | ipsource | iptos | ipttl | macdest |
macsource | port | tcpsourceportIP | udpsourceportIP | tcpdestportIP |
udpdestportIP] [all-traffic-entries |data ] {maskmask } {port-string
port_string}]

```


NetLogin Commands

```

configure netlogin idle-timeout {dot1x | mac | web-based} timeout

configure netlogin session-timeout {dot1x | mac | webbased} timeout

configure netlogin ports [all | port_list] [allowedusersallowed_users |
authentication mode [optional | required] | trap [all-traps | no-traps |
[ {success} {failed} {terminated} {max-reached} ]]]

configure netlogin trap max-users [enable | disable]

show netlogin timeout

show netlogin session {all | summary} {mac-address mac_address} {portsports}
{agent [dot1x | mac | webbased]}

show netlogin trap

unconfigure netlogin [idle-timeout | session-timeout] {dot1x | mac | web-based}

unconfigure netlogin ports [all | port_list] [allowedusers | authentication mode]

clear netlogin session [ {mac mac_address} {port ports} ] {dot1x | mac | web-based}

```

Access Control List (ACL) Library Enhancements

To implement ONEPolicy requires enhancements to certain existing Access Control List (ACL) conditions and actions, plus the addition of some new ones:

- **ttl**—New condition with optional mask. Matches IPv4 Time-To-Live and IPv6 Hop Limit.

Syntax: `ttl value {mask value}`

- **add-vlan-id**—New action that adds a new outer VLAN ID. If the packet is untagged it adds a VLAN tag to the packet. If the packet is tagged, it adds an additional VLAN tag. Only supported in VLAN lookup stage (VFP).

Syntax: `add-vlan-id value`

- **ip-tos**—Existing condition that now accepts an optional mask.

Syntax: `ip-tos value {mask value}`

- **vlan-format**—New condition matches packets based on its VLAN format. Can be one of the four values:
 - **untagged**—all untagged packets
 - **single-tagged**—all packets with only single tag
 - **double-tagged**—all packets with double tag
 - **outer-tagged**—all packets with at least one tag (single tag or double tag)

Syntax: `vlan-format untagged | single-tagged | doubletagged | outer-tagged`

- **replace-dscp-value**—New action that replaces the existing Differentiated services code point (DSCP) value of the packet.

Syntax: `replace-dscp-value value`

- **ethernet-type**—Existing condition that can now accept an optional mask.

Syntax: `ethernet-type value {mask value}`

- **destination-port**—Existing condition that can now accept an optional mask.

Syntax: `destination-port value {mask value}`

- **source-port**—Existing condition that can now accept an optional mask

Syntax: `source-port value {mask value}`

- **fragments**—Existing condition that matches any fragment of a fragmented packet, including the first fragment. Previously, this condition didn't include the first fragment.

Syntax: `fragments`

- **first-fragments**—Existing condition that matches only the first fragment of a fragmented packet. Previously, this condition also matched non-fragmented packets.

Syntax: `first-fragments`

- **do-ipfix**—New action that records the matching packet. Can be used on both ingress and egress.

Syntax: `do-ipfix`

- **do-not-ipfix**—New action that cancels recording for the matching packet. Can be used to reduce demand on egress IPFIX capacity (and to reduce recording loss) during packet flooding situations. For example, egress ACLs that recognize broadcast and/or IP multicast packets could prevent egress IPFIX recording. Can be used on both ingress and egress.

Syntax: `do-not-ipfix`

- **redirect-port-copy-cpu-allowed**—The existing "redirect-port" action in ACL includes three hardware actions: RedirectPort, CopyToCpuCancel and GpDropCancel. Because of this, the "redirect-port" action cannot be used with another actions that try to copy a packet to CPU, like "copy-cpu-sdn". This new action redirects a packet out of an output port, but does not enforce a requirement that Copy to CPU must be canceled.

Syntax: `redirect-port-copy-cpu-allowed value`

- **redirect-port-list-copy-cpu-allowed**—Same as `redirect-port-copy-cpu-allowed`, but allows redirect to a list of ports.

Syntax: `redirect-port-list-copy-cpu-allowed value {,value}`

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- **vlan-format** mistakenly identifies untagged packets as tagged in the IFP stage for the following switches: Summit X480, Summit X650, BlackDiamond 8900-G96T-c, BlackDiamond 8900-10G24X-c, BlackDiamond 8900-G48T-xl, BlackDiamond 8900-G48X-xl, and BlackDiamond 8900-10G8X-xl.
- **fragments** is partially supported on the BlackDiamond G48Te2 I/O modules. On this modules, this condition only matches fragmented packets and the last fragmented packet, and does not match the first fragment of the packet.
- **add-vlan-id** is only available on switches with VFP stages.
- IPFIX actions are only supported on Summit X460, X460-G2, and X480 series switches, BlackDiamond 8900-xl and -96T modules, and BlackDiamond X8-100G4X and BDX X8 xl-series modules.

Class of Service (CoS)

Class of Service (CoS) prioritizes, rate-limits, rate shapes, and otherwise controls defined traffic types of a switch; it is used as part of a bandwidth management strategy. CoS is an enhancement to the existing Quality of Service (QoS) feature. CoS is typically configured using NetSight through the CoS Management Information Base (MIB).

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- CoS reference tables have a fixed mapping.
- Both global and per-port meters cannot be configured at the same time on platforms that require the port meter map table.
- Per-port meter out-of-profile counters are not accurately supported.
- Flood out-of-profile counters and actions are the aggregate of all flood types (unknown unicast, multicast, and broadcast) for a given port.
- On Summit X670-G2 and X770 series switches and on BlackDiamond X8 modules: BDXB-100G4X, BDXB-100G4X-XL, and BDXB-40G12X-XL, when hybrid scheduling is configured on ports (WxRR with some queues in Strict-Priority), the SP queues must be contiguous

New CLI Commands

```
create ports group port_group
```

```
delete ports group port_group
```

```
configure ports {group} port_group [[add | delete] port_list]
```

```
show ports group {port_group}
```

```
unconfigure {meter} metername ports [port_group | port_list]
```

```

clear ports [all | port_list|port_group] rate-limit flood out-of-profile
{disabled-ports} {status | counter}

unconfigure cos-index cos_index [{qosprofile} {ingress-meter} {replace-tos}]

show cos-index {cos_index} show meter {metername} out-of-profile {{disabled-ports} ports [port_list | port_group] | global-count}

clear meter {metername} out-of-profile {disabled-ports} {status | counters}
{ports [ all | port_list | port_group]}

configure ports [port_list | all] dot1p dot1p_priority

unconfigure qoscheduler ports [port_list | port_group | all]

```

Changed CLI Commands

Changes are underlined.

```

configure qoscheduler [strict-priority | weighted-round-robin | weighted-deficit-round-robin] {ports [port_list | port_group | all]}

show qoscheduler {ports [port_list | port_group | all]}

configure {qosprofile} qosprofile [{maxbuffer buffer_percentage} {weight weight_value | use-strict-priority} {ports [port_list | port_group | all]}]

configure [{qosprofile} {egress} qosprofile | {qosprofile} ingress [iqosprofile]]
[{minbw minbw_number} {qos-weight weight} {maxbw maxbw_number} | {committed_rate committed_bps [K | M]} {qos-weight weight} {peak_rate peak_bps [K | M]} ]
{priority [priority | priority_number]} [ports [port_list | port_group | all]]

show qosprofile {ingress | egress} [ports [port_list | all | port_group] | NULL]

configure {qosprofile} {egress} qosprofile [wred [color [tcp [green | red] | non-tcp [any | red]]] {min-threshold min_thresh} {max-threshold max_thresh} {max-drop-rate max_drop_rate} | avg-weight avg_weight]] [ports [port_list | port_group | all]]

configure {meter} metername [{committed-rate cir committed-rate-unit {committed-burst-size committed-burst-size [Kb|Mb]}}] {peak-rate pr peak-rate-unit} {peak-burst-size peak-burst-size [Kb|Mb]}] {max-burst-size burst-size [Kb | Mb]}}] {out-actions [{disable-port} {drop | set-drop-precedence {dscp [dscp-value | none]}}] {log} {trap}}] {ports [port_group | port_list]}

show meter {meter_name} {ports [port_group|port_list]} {out-actions}

configure ports [port_list|port_group] rate-limit [egress [no-limit | cir-rate [Kbps | Mbps | Gbps] {max-burst-size burst-size [Kb | Mb]}}] | flood [broadcast | multicast | unknown-destmac] [no-limit | pps {out-actions [{log} {trap} {disable-port}}]]]

```

```

configure dot1p type dot1p_priority {qosprofile} qosprofile {ingress-meter
[ing_meter | none]}

[enable | disable] dot1p replacement ports [port_list | all] {{qosprofile}
qosprofile}

[enable | disable] diffserv replacement ports [port_list | all] {{qosprofile}
qosprofile}

show ports {port_list | port_group | tag tag} rate-limit flood {out-actions |
out-of-profile {disabled-ports}} {no-refresh}

```

Note



For the above command, the show screen for the rate-limit flood configuration now shows the **port-group** option. A separate screen shows the configured out-actions. Another screen shows a per-port out-of-profile status that indicates that the flood limit is exceeded on the port. Another screen show ports that are disabled.

The `show access-list meter` command now shows per-port meter out-of-profile counters for meters that are applied using ACL rules. The output shows the additional syslog,

The `clear access-list meter` command now allows the clearing of the out-of-profile per-port meters counters. When a per-port meter is specified, it clears the counter for the rule associated with the meter.

The `show port information details` command displays the default dot1p priority used for the internal priority for untagged traffic on a specified port. Additionally, it displays the per-port/per-qosprofile dot1p and diffserv examination status.

Command Usability Enhancements

This feature changes a select set of commands so that you may specify VLANs by VID instead of by name. Some commands allow you to specify a list of VIDs.

Additionally, ExtremeXOS 16.1 introduces a new show system command that aggregates the output of the following commands:

- `show switch`
- `show version`
- `show temperature`
- `show power`
- `show fans`
- `show odometers`

Additionally, there is a new `cli refresh` command to set the auto refresh behavior for certain show commands. The `cli refresh` command controls the default behavior of show commands that have no-refresh/refresh options if neither is specified. Previously, these show commands had only a **no-refresh** option, but now have both so that the global setting can be overridden.

Finally, there is a new command, `configure cli [{lines height} {columns width}]`, that configures the number of lines and columns for the current logon session only.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- For commands that accept a list of VLANs, if some VLANs in the list are unresolvable, the command continues to execute for the remaining VLANs.
- The command show configuration shows the individual commands with the VLAN names, not VLANs or VLAN lists.

New CLI Commands

```
show system [enable | disable]
```

```
cli refresh {session | permanent}
```

```
configure cli [{lines height} {columns width}]
```

Changed CLI Commands

The following commands are modified to allow you to specify VLANs by VID (or a list of VLANs) instead of by name:

```
configure mirror {mirror_name} add [vlanvlan_id {ingress | [portport {ingress}]} |
portport vlan vlan_id {ingress}]
```

```
configure mirror {mirror_name} delete [vlanvlan_id {port port} | port port
vlanvlan_id]
```

```
configure vlan vlan_id add secondary-ipaddress [ipaddress {netmask} | ipNetmask]
```

```
configure vlan vlan_id delete secondary-ipaddress [all | ipaddress]
```

```
configure vlan vlan_id ipaddress [ipaddress {netmask} | ipNetmask]
```

```
configure vlan vlan_id name new_name
```

```
clear l2stats vlan vlan_list
```

```
configure ip-mtu mtu vlan vlan_list
```

```
configure ports [port_list | all] monitor vlanvlan_list {rx-only | tx-only}
```

```
configure ports port_list {tagged} vlan vlan_list [limit-learning number {action
[blackhole | stop-learning]} | unlimited-learning]
```

```
configure ports port_list {tagged} vlan vlan_list [lock-learning | unlock-
learning]
```

```
configure vlan vlan_list add ports [port_list | all] {tagged | untagged |
private-vlan translated}
```

```

configure vlan vlan_list delete ports [port_list | all]

configure vlan vlan_list protocol {filter} filter_name

configure vlan vlan_list {qosprofile} [qosprofile | none]

create vlan vlan_list {vr vr-name} {description vlan-desc}

delete vlan vlan_list

[enable | disable] iparp gratuitous protect vlan vlan_list

[enable | disable] ipforwarding {ipv4 | ipv6} vlan vlan_list

[enable | disable] learning vlan vlan_list

[enable | disable] loopback-mode vlan vlan_list

[enable | disable] vlan vlan_list

unconfigure vlan vlan_list ipaddress

configure vman vman_id add ports [port_list | all] {tagged | untagged {port-cvid
port_cvid} | cep cvid cvid_first {- cvid_last} {translate cvid_first_xlate {-
cvid_last_xlate}}}

configure vman vman_id ipaddress [ipaddress {netmask} | ipNetmask]

configure vman vman_id ports [port_list | all] add cvid cvid_first {- cvid_last}
{translate cvid_first_xlate {- cvid_last_xlate}}

configure vman vman_id name new_name

configure vman vman_list delete ports [port_list | all]

configure vman vman_list ports [port_list | all] delete cvid cvid_first {-
cvid_last}

configure vman vman_list protocol {filter} filter_name

configure vman vman_list {qosprofile} [qosprofile | none]

create vman vman_list {vr vr-name} {description vlan-desc}

delete vman vman_list

[enable | disable] learning vman vman_list

unconfigure vman vman_list ipaddress

show ip dad vlan vlan_list {tentative | valid | duplicate}

show ipv6 dad vlan vlan_list {tentative | valid | duplicate} {detail}

show l2stats vlan vlan_list

```

```

show vlan vlan_list statistics

show [vlan vlan_list | vman vman_list] {ipv4 | ipv6}

configure vlan vlan_list add ports [all | port_list] {tagged | untagged} {stp}
stp_name {dot1d | emistp | pvst-plus}

configure {stp} stp_name add vlan vlan_list ports [all | port_list] {dot1d |
emistp | pvst-plus}

configure {stp} stp_name delete vlan vlan_list ports [all | port_list]

[enable | disable] {stp} stp_name auto-bind vlan vlan_list

show vlan vlan_list stp {blocked-ports}

show fdb vlan vlan_list {netlogin [all | mac-based-vlans]}}

show fdb stats vlan vlan_list

show iparp security vlan vlan_list

show iparp stats vlan vlan_list

show iparp vlan vlan_list

show neighbor-discovery {cache {ipv6}} vlan vlan_list

show vlan vlan_list security

show netlogin {port port_list} vlan vlan_list

show netlogin {port port_list} vlan vlan_list dot1x detail

show netlogin guest-vlan vlan_list

show netlogin authentication [service-unavailable | failure] vlan vlan_list

show ip-security arp learning vlan vlan_list

show ip-security arp validation vlan vlan_list

show ip-security arp validation violations vlan vlan_list ports [ports | all]

show ip-security dhcp-snooping vlan vlan_list

show ip-security dhcp-snooping entries vlan vlan_list

show ip-security dhcp-snooping information-option circuit-id vlan-information
vlan vlan_list

show ip-security dhcp-snooping violations vlan vlan_list

```

The following show commands are modified so that you can now, per command, set the auto-refresh behavior:


```

debug hal show ports {port_list} qosmonitor {congestion} {no-refresh | refresh}

show fdb stats [ports {all | port_list} | vlan {all} | {vlan} vlan_name ] {no-
refresh | refresh}

show iparp stats [[ vr_name | vr {all | vr_name} ] {no-refresh | refresh} | {vr}
summary ]

show iparp stats [vlan {all {vr vr_name}} | {vlan} vlan_name] {no-refresh |
refresh}

show iparp stats ports {all | port_list} {no-refresh | refresh}

show fdb mac-tracking statistics {mac_addr} {no-refresh | refresh}

show trill ports {port_list} [counters {no-refresh | refresh | detail}]

show ports {port_list | tag tag} collisions {no-refresh | refresh} {port-number}

show ports {port_list | tag tag} configuration {no-refresh | refresh} {port-
number}

show ports {port_list | tag tag | stack-ports {tack_port_list}} txerrors {no-
refresh | refresh} {port-number}

show ports {port_list | tag tag} packet {no-refresh | refresh} {port-number}

show ports {port_list | tag tag} wan-phy errors {no-refresh | refresh}

show ports {port_list | tag tag} wan-phy events {no-refresh | refresh}

show ports {port_list} wan-phy overhead {no-refresh | refresh} {port-number}

show ports {port_list | tag tag | stack-ports {stack_port_list}} rxerrors {no-
refresh | refresh} {port-number}

show ports {port_list | tag tag} anomaly {no-refresh | refresh} {port-number}

show ports {port_list | port_group | tag tag} rate-limit flood {out-actions |
out-of-profile {disabled-ports}} {no-refresh | refresh} {port-number}

show ports {port_list} ip-fix {no-refresh | refresh} {port-number}

show ports {port_list | tag tag | stack-ports {stack_port_list}} statistics {no-
refresh | refresh} {wide} {port-number}

show ports {port_list | tag tag} qosmonitor {egress | ingress} {congestion}
{packets | bytes} {no-refresh | refresh} {port-number}

show ports {port_list | tag tag} flow-control {rx-pauses | tx-pauses} {no-refresh
| refresh} {port-number}

show ports {port_list} wred {no-refresh | refresh} {port-number}

```

```

show ports {port_list} eee {no-refresh | refresh} {port-number}

show ports {port_list | tag tag} congestion {no-refresh | refresh} {port-number}

show ports {port_list} vlan statistics {no-refresh | refresh} {port-number}

show [vlan | {vlan} vlan_name] statistics {no-refresh | refresh}

show ports {port_list | tag tag} {no-refresh | refresh}

show ports {port_list} tdm errors {near-end} {total | intervals | current {no-
refresh | refresh}}

show ports {port_list} tdm configuration {no-refresh | refresh} {port-number}

show ports {port_list} tdm {no-refresh | refresh}

show ports {port_list} tdm alarms {no-refresh | refresh}

show ports {port_list} dot1p out-of-profile {disabled-ports} {no-refresh |
refresh}

show ces {ces_name} errors {total | intervals | dayIntervals | current {no-
refresh | refresh}}

```

Access Control List (ACL) Two-Stage Policy

This feature exposes the VLAN Content Aware Processor/VLAN Filter Processor (VCAP/VFP) using the ExtremeXOS Access Control List (ACL) manager.

The VCAP/VFP is used to filter packets before ingress processing. It can be used to assign the VLAN, set a class ID, or perform other more traditional ACL actions, such as drop or count. In general, this stage's scale, actions, and match criteria are more limited than the ingress stage.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, and X450-G2 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

- The VFP match criteria, scale, and actions are more limited than that of the regular ingress ACLs—Ingress Content Aware Processor/Ingress Filter Processor (ICAP/IFP).
- Rule actions in the VFP can be overridden by rule actions in the IFP.
- Packets are always presented to the IFP even when the VFP drops the packet
- The 'vlan-id' match criteria only works on packets received with an 802.1Q tag in the packet.

New CLI Commands

New ACL action modifier:

```
class-id value 0-4095
```

This action can be specified on any rule within a policy file or within a list of dynamic access-lists. When specified, this action signifies that the rule is installed in the “LOOKUP stage” access-list resource (VFP).

New ACL match criteria:

```
class-id value 0-4095
```

This match condition can be specified on any rule within a policy file or within a list of dynamic access-lists. A rule cannot both match a class-id and specify a class-id as an action. When a “class-id” match criteria is specified, the associated rule is programmed into the normal “INGRESS stage” access-list hardware resource (IFP).

Security Enhancements

This feature includes the following changes and enhancements:

- Configurable timed lockout that is applied to accounts after a configurable number of failed logon attempts.
- Stronger hash algorithm for account passwords.

Note



Due to the stronger hash algorithm, if you create accounts in , and then downgrade to versions earlier than , you may encounter problems using the passwords for these accounts. For more information about this issue, visit:

<http://extr.co/1KfSszY>

- Removal of unmasked passwords in the command line interface.
- Stronger obfuscation of RADIUS and TACACS+ shared secrets.
- Integrity checking of downloaded images.
- Syslog alert issued when a configurable percentage of the Syslog memory buffer is filled.
- Optionally restricting the use of `show log` and `show diagnostics` commands by non-administrator accounts.
- The “safe defaults” script (unconfigured switch startup wizard) enables these new options collectively, as well as forcing the user to change the default administrator and failsafe passwords.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches
- E4G-200 and E4G-400 cell site routers

New CLI Commands

```
configure account [all | <name>] password-policy lockout-time-period [num_mins | until-cleared]
```

```
configure log target memory-buffer alert percent-full [percent | none]
```

```
configure cli password prompting-only [on | off] configure log messages privilege [admin | user]
```

```
configure diagnostics privilege [admin | user]
```

Changed CLI Commands

The output of the this command now displays account lockout time period information:

```
show accounts password-policy
```

If a downloaded image does not have a signature, a warning message appears. You may choose to continue or terminate the installation:

```
download image [[hostname | ipaddress] filename] {{vr} vrname} {block-size
block_size} | memorycard filename] {partition} {slot slot number
```

The log buffer percentage full and configurable percentage threshold information appears in the output of the following command:

```
show log configuration {target {upm {upm_profile_name} | xml-notification
{xml_target_name} | console | session | memory-buffer | primary-msm | primary-mm
| primary-node | backup-msm | backup-mm | backup-node | nvram | syslog
{ipaddress|ipPort} {vr vr_name} {local}} | filter {filter-name}}
```

The following command shows the current password prompting setting:

```
show management
```

Generalized Precision Time Protocol (gPTP) Enhancement

Previously, the number of Generalized Precision Time Protocol (gPTP)-capable ports was static. Switches now handle a variable number of ports based on the number of ports on the switch.

Limitations

- ExtremeXOS 16.1 slave ports synchronize to grandmasters, such as Symmetricom, and to other ExtremeXOS 16.1 clocks, but not to ExtrememXOS 15.7, and earlier. If networks of clocks are to be upgraded to ExtremeXOS 16.1, complete the upgrades simultaneously or staged starting closest to the grandmaster. Before beginning a staged upgrade, where an earlier version of ExtremeXOS must synchronize to an ExtremeXOS 16.1 clock, test the particular configuration beforehand.
- ExtremeXOS 16.1 slave clock ports must be configured with the “slave-only” option to synchronize to other ExtremeXOS 16.1 clocks.

Extreme Loop Recovery Protocol (ELRP) Port Shutdown

Extreme Loop Recovery Protocol (ELRP) detects loops by sending out an ELRP protocol data units (PDUs) out of one or more ports of a particular VLAN. ELRP takes system MAC addresses, changes them to a broadcast MAC address by appending “01:” to the front, and then sends out the PDUs. If PDUs are received back by ELRP, a loop is present. Each ELRP PDU is sent on a particular VLAN, so you must configure each VLAN that you wish to monitor.

In this ELRP enhancement, when a loop is detected using ELRP, an option to disable the port where the ELRP packet egresses is added to suppress the loop. You may specify a duration, which after it expires, the ports are enabled, or you can keep the ports disabled permanently until you choose to enable them.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, and X450-G2 series switches
- E4G-200 and E4G-400 cell site routers

Limitations

Entire port is blocked regardless of which VLAN that the loop was detected on.

Changed CLI Commands

The last six options in the following command were changed:

```
configure elrp-client periodic vlan_name ports [ports | all] interval sec [log |
log-and-trap | trap] {disable-port {egress | ingress} {duration {seconds} |
permanent }}
```

The output of the `show elrp` command now shows egress/ingress information in the “Action” column and there is a new column (Disable Port) showing disabled port status.

The output of the `show elrp disabled-ports` command now shows egress/ingress information in the new “Disable Direction” column.

Increase of Protocol-Independent Multicast (PIM) Control Packets

Previously, ExtremeXOS Protocol-Independent Multicast (PIM) implementation sends maximum 1,500 byte-size control packets that can accommodate 175 sources per multicast group.

The size of these control packets has now been increased to accommodate 3,000+ sources per group. This large control packet is fragmented at the IP layer and reassembled at the received node.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, and X450-G2 series switches
- E4G-200 and E4G-400 cell site routers

OpenFlow Updated Match Conditions and Actions

ExtremeXOS 16.1 includes an upgrade to OpenFlow v1.3 by upgrading from version 1.4 to 2.1 of OpenVswitch. The match conditions and actions tables change as result of this upgrade.

Supported Platforms

- BlackDiamond X8 [all modules; single Management Module (MM) only]
- BlackDiamond 8800 [8900 (XL-series) and C-series; single Master Switch Fabric Module (MSM) only]

- Summit X770, X670, X480, X460, and X440

Alternate Stacking Supported on 1G Variant of Summit X460-G2 Series Switches

Alternate stacking is now available for 1G oriented Summit X460-G2 series switches with 10G VIM modules (VIM-2T or VIM-2X).

Supported Platforms

- Summit X460-G2-24t-GE4
- Summit X460-G2-24p-GE4
- Summit X460-G2-48t-GE4
- Summit X460-G2-48p-GE4

Two-Way Active Measurement Protocol (TWAMP) Light

This feature is the light version of TWAMP, which is an industry standard (RFC 5357) for measuring round-trip performance between two devices that support the TWAMP protocols. TWAMP defines two protocols: the TWAMP-Control protocol and the TWAMP-Test protocol. The TWAMP-Control protocol is used to set up test sessions. The test sessions use the TWAMP-Test protocol to transmit and reflect performance measurement packets. The TWAMP-Control protocol uses TCP for communication, while the TWAMP-Test protocol uses UDP.

TWAMP defines four logical roles; Session-Sender, Session-Reflector, Server, and Control-Client. These logical roles can be split between two entities to form a client/server paradigm (referred to as the two-host implementation in the RFC). The logical roles of the Control-Client and the Server communicate using the TWAMP-Control protocol to set up test sessions. Each test session consists of a transmitter and a reflector, fulfilled by the logical roles of Session-Sender and Session-Reflector respectively.

The Control-Client initiates a TCP connection to the well-known TWAMP-Control port 862. After the TCP connection is established, the Server transmits the first message by sending the 'Server Greetings' message. The Control-Client responds with a 'Setup Response' message. The three-way TWAMP-Control handshake is completed when the Server responds with the 'Server Start' message. The 'Control-Client' is now able to transmit command messages to the Server.

The test sessions are set up using the 'Request-Session' command message, sent from the Control-Client to the Server. The Server replies with an 'Accept Session' message, which indicates if the Server is capable of accommodating the request. The Control-Client may send several 'Request-Session' command messages to set up multiple test sessions. To begin the tests, the Control-Client transmits a 'Start-Session' command message. The Server replies with a 'Start-Ack' message. The Control-Client does not begin its test until it receives the 'Start-Ack' message. This allows the Server ample time to configure the test sessions. The Control-Client stops the test sessions with a 'Stop-Sessions' message. The Server does not respond to this message.

This feature, TWAMP Light, consists of the TWAMP logical role of the Session-Reflector as defined in Appendix I of RFC 5357. This light version of TWAMP contains two entities: the Client entity takes on the TWAMP logical roles of Session-Sender, Server, and Control-Client, while the Server takes on the TWAMP logical role of the Session-Reflector. To establish TWAMP Light, you must configure endpoints,

which define the destination of TWAMP-Test packets generated by the Client. An endpoint receiving a new TWAMP-Test packet creates a test session consisting of the following four-part tuple; client IP address, client UDP port, endpoint IP address, and endpoint UDP port. The tuple does not include the VR because it requires the Default VR for the first phase. A session timeout value, configured globally, determines the amount of time test sessions exist after the last reception of a TWAMP-Test packet. Test sessions are used to keep track of the session data, such as the sequence number. The Session-Reflector still responds to TWAMP-Test packets that do not match an existing test session or if a new test session cannot be created due to lack of resources.

Platforms Supported

- Summit X440, X460
- BlackDiamond 8800 with MSM128, MSM48c, and MSM-96

Limitations

- Endpoints may only be created on the default virtual router (VR-Default)
- Limit of 256 endpoints Maximum of 2,000 test sessions
- TWAMP-Test error-estimate field does not reflect Network Time Protocol (NTP) or Simple Network Time Protocol (SNTP) timing values

New CLI Commands

```
[enable | disable] twamp reflector
```

This command enables or disables the Session-Reflector. If you disable the Session-Reflector, the application terminates all current TWAMP test sessions.

```
configure twamp reflector sessions count timeout ref_wait
```

This command allows you to modify the number of test sessions to support and timeout value for those test sessions. The timeout value is the REFWAIT value specified in RFC 5357.

<count>: range 0 - 2000 entries; default 2000

<ref_wait>: range 30 - 3600 seconds; default 900 seconds

```
(un)configure twamp [add|delete] endpoint {vr name} ipaddress ip port udp_port
```

This command allows you to configure (and unconfigure using `unconfigure`) the TWAMP endpoints. You specify the IP address and UDP port number for the endpoint. Removing the endpoint terminates all test sessions associated with the endpoint.

<ip>: The endpoint IP address, either IPv4 or IPv6

<udp_port>: The UDP port the endpoint listens on; range is 1,025-65,535

<name>: An optional VR may be used; default is VR-Default

```
show twamp reflector
```

This command displays the configured values and runtime information of the Session-Reflector and its endpoints:

```
Session Information
Used: 165 of 200
Timeout: 300 seconds
```

| Endpoints | Port | Sessions | Rx Packets | Tx Packets |
|------------|------|----------|------------|------------|
| 19.1.1.100 | 5000 | 5 | 3091 | 3091 |
| 19.1.1.100 | 5001 | 40 | 5521 | 5521 |
| 19.1.1.100 | 5002 | 40 | 4728 | 4728 |
| 19.1.1.100 | 5003 | 40 | 3916 | 3916 |
| 18.1.1.100 | 5000 | 40 | 9266 | 9266 |

Displayed 5 endpoints

```
show twamp endpoint ipaddress ip port udp_port
```

This command displays the endpoint configured values, runtime data, and test session information. Specifying an IP address and port are optional:

```
TWAMP Endpoint
Endpoint Information
  Local Address: 19.1.1.100      Listening Port:          5000
  Received Packets: 7948        Transmitted Packets: 7948
  Active Sessions: 5

Session created on Thu Nov 13 15:41:49 2014
  Peer Address: 19.1.1.2        Port: 11001
  Sequence Number: 1575         Last recv'd packet: 84ms

Session created on Thu Nov 13 15:41:49 2014
  Peer Address: 19.1.1.2        Port: 11002
  Sequence Number: 1555         Last recv'd packet: 16ms

Session created on Thu Nov 13 15:41:49 2014
  Peer Address: 19.1.1.2        Port: 11003
  Sequence Number: 1595         Last recv'd packet: 241ms

Session created on Thu Nov 13 15:41:49 2014
  Peer Address: 19.1.1.2        Port: 11004
  Sequence Number: 1729         Last recv'd packet: 249ms

Session created on Thu Nov 13 15:41:49 2014
  Peer Address: 19.1.1.2        Port: 11005
  Sequence Number: 1489         Last recv'd packet: 15ms

Displayed 5 sessions
```

Flow Redirects (Policy-Based Routes) Limits Increase

The limit of number of flow redirects and flow redirect next hops has been increased. This allows you to install many more policy-based routes in the switch.

The number of flow redirects was limited to 256; that limit is increased to 4,096. The number of next hops was limited to 32 per flow redirect; this limit is increased to 4,096 next hops cumulatively across all flow redirects.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches

- E4G-200 and E4G-400 cell site routers

Single Virtual Group for User Access Control Lists (ACLs)

This feature allows you to put all user rules into a single virtual group to prevent multiple rule matches and allow only the highest priority rule to do the matching and execute its actions.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670, X670-G2, X480, X460, X460-G2, X450-G2, X440, and X430 series switches
- E4G-200 and E4G-400 cell site routers

40Gbps LR4 Parallel Single-Mode (PSM) Quad Small Form-Factor Pluggable (QSFP) Optical Transceiver

This feature adds support for the LR4 Parallel Single Mode (PSM) Quad Small Form-Factor Pluggable (QSFP) optical transceiver on 40G optical QSFP+ ports. Running in 4 × 10g mode allows gives you the capability of having four independent transmit and receive channels, each capable of 10Gbps operation over a 10km single mode fiber hydra MPO to 4xLC duplex patch cord terminated with standard 10G LR SFP+ optical transceivers.

Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800 series switches
- Summit X770, X670-G2, and X480series switches
- Summit X670-G2-48x (stacked)

Changed CLI Commands

The following commands' outputs now display the PSM QSFP media type when detected:

```
show port configuration
```

```
show port transceiver information detail
```

ExtremeXOS Images for Summit X480 Series Switches

Due to additional functionality and new platforms supported, the ExtremeXOS 15.6 and later software image is too large to download onto the Summit X480 series switches. To resolve this issue, Summit X480 series switches now have two separate software image files used for both individual switches and stacks that include Summit X480 series switches.

Table 3: Summit X480 Series Switches Software Image Files

| | Main Install image | Diagnostic image |
|-----------|--|--|
| Content | All Summit X480 content (except diagnostics) | Summit X480 diagnostics |
| File Name | summitX480-16.1.xx.yy.xos | summitX480-16.1.xx.yy-diagnostics.xmod |

Table 3: Summit X480 Series Switches Software Image Files (continued)

| | Main Install image | Diagnostic image |
|--------------------|--|---|
| File Type | Standard ExtremeXOS image | XMOD image |
| Installation Notes | <ul style="list-style-type: none"> Installing the main SummitX480 image over a previous release leaves the previous installation of the diagnostics image intact, as it is stored separately from the main ExtremeXOS image. You can continue to use the previously installed diagnostic version to run diagnostics. The Summit XMODs, such as SSH can be used with the summitX480 ExtremeXOS image. | To update to a newer version of the diagnostics, you download and install the latest XMOD version. The diagnostics XMOD can be installed to the active or standby partition and diagnostics can be used immediately. There is no need to reboot or any other action to complete the installation. |

The following scenarios will produce an error or warning message:

- Not having the diagnostic image installed on a Summit X480 series switch or slot.
- Installing the main Summit X480 image without the diagnostics image present.
- Installing the general Summit image (summitX-16.1.xx.yy.xos, rather than the Summit X480-specific image) on a Summit X480 series switch.

**Note**

If Summit X480 series switches require rescue recovery, you can use the `summitX-16.1.xx.yy.xos` file image, and this image installs the diagnostics capability.

New Hardware Supported in 16.1.3

This section lists the new hardware supported in ExtremeXOS 16.1.3:

- BlackDiamond 8800-MSM96 Management Switch Fabric Module

New Hardware Supported in ExtremeXOS 16.1

This section lists the new hardware supported in ExtremeXOS 16.1:

- BDXB-40G12X-XL I/O, BDXA-G48T, and BDXA-G48X modules for the BlackDiamond X8 series switches
- Summit X450-G2 series switches:

24t-10GE4, 24p-10GE4, 48t-10GE4, 48p-10GE4, 24t-GE4, 24p-GE4, 48t-GE4, 48p-GE4

CLI Command Output Format of Ports Lists

For ExtremeXOS 16.1 and later, the output of CLI commands showing ports lists does not display spaces between commas.

For example: “3:1,7:13” instead of “3:1, 7:13”

Joint Interoperability Test Command (JITC) Compliance

If you require Joint Interoperability Test Command (JITC) compliance, you can use the command `configure snmp compatibility get-bulk reply-too-big-action [standard | too-big-error]` to change ExtremeXOS from Ridgeline-compatible mode (standard), the default mode, to JITC-compliant mode (too-big-error).

Please note that switching to JITC-compliant mode causes Ridgeline to display potentially unreliable information.

Extreme Hardware/Software Compatibility and Recommendation Matrices

The *Extreme Hardware/Software Compatibility and Recommendation Matrices* provide information about the minimum version of ExtremeXOS software required to support switches, as well as SFPs, XENPAKs, XFPs, and other pluggable interfaces.

This guide also provides information about which optics are supported on which hardware platforms, and the minimum software version required.

The latest version of this and other ExtremeXOS guides are at: <http://documentation.extremenetworks.com>

Compatibility with NetSight

ExtremeXOS 16.1.3 is compatible with NetSight version 6.2.0.220 and later.

Upgrading to ExtremeXOS

For instructions about upgrading ExtremeXOS software, see "Software Upgrade and Boot Options" in the *ExtremeXOS User Guide*. The following are miscellaneous hitless upgrade notes:

- Beginning with ExtremeXOS 12.1, an ExtremeXOS 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 error message `Error: Image can only be installed to the non-active partition.` appears. An ExtremeXOS modular software package (.xmod file) can still be downloaded and installed on either the active or alternate partition.
- SummitX software is required for E4G cell site routers.
- Beginning with ExtremeXOS 15.4, a limited hitless upgrade procedure is supported on the BlackDiamond X8 and BlackDiamond 8800 series switches.
- For Summit X480 series switches, starting with ExtremeXOS 15.6, two separate software image files are used for both individual switches and stacks that include Summit X480 series switches. For more information, see [ExtremeXOS Images for Summit X480 Series Switches](#).

Downloading Supported MIBs

The Extreme Networks MIBs are located on the eSupport website under **Download Software Updates**, located at: <https://esupport.extremenetworks.com>.

You need to provide your serial number or agreement number, and then the MIBs are available under each release.

Tested Third-Party Products

This section lists the third-party products tested for 21.1.

Tested RADIUS Servers

The following RADIUS servers are fully tested:

- Microsoft—Internet Authentication Server
- Meetinghouse
- FreeRADIUS

Tested Third-Party Clients

The following third-party clients are fully tested:

- Windows 7
- Windows Vista
- Linux (IPv4 and IPv6)
- Windows XP (IPv4)

PoE Capable VoIP Phones

The following PoE capable VoIP phones are fully tested:

- Avaya 4620
- Avaya 4620SW IP telephone
- Avaya 9620
- Avaya 4602
- Avaya 9630
- Avaya 4621SW
- Avaya 4610
- Avaya 1616
- Avaya one-X
- Cisco 7970
- Cisco 7910
- Cisco 7960
- ShoreTel ShorePhone IP 212k
- ShoreTel ShorePhone IP 560
- ShoreTel ShorePhone IP 560g
- ShoreTel ShorePhone IP 8000
- ShoreTel ShorePhone IP BB 24
- Siemens OptiPoint 410 standard-2
- Siemens OpenStage 20

- Siemens OpenStage 40
- Siemens OpenStage 60
- Siemens OpenStage 80

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

Service Notifications

To receive proactive service notification about newly released software or technical service communications (for example, field notices, product change notices, etc.), please register at: www.extremenetworks.com/support/service-notification-form

2 Limits

This chapter summarizes the supported limits in ExtremeXOS 16.1.3-Patch1-7.

[Table 4: Supported Limits](#) on page 47 summarizes 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 ExtremeXOS books.



Note

The term “BlackDiamond 8000 e-series” refers to all BlackDiamond 8500 e-series and 8800 e-series modules. The term “BlackDiamond 8000 series” refers to all BlackDiamond 8500, 8800, and 8900 series modules.

The scaling and performance information shown in [Table 4](#) 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 [Table 4](#) 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 table.

On products other than the BlackDiamond 8900 xl-series, BlackDiamond X8 series, and Summit X480 series, it is not advised to have greater than 25,000 total IP routes from all routing protocols. Adverse effects can occur with routing tables larger than this, especially when a single network event or CLI command affects a significant number of routes. For example, just after such a network event, the added system load will cause a save configuration command to time out.

Table 4: Supported Limits

| Metric | Product | Limit |
|--|--|---|
| AAA (local) —maximum number of admin and local user accounts. | All platforms | 8 |
| Access lists (meters) —maximum number of meters. | BlackDiamond 8000 series e-series, group of 24 ports c-series | 512 2,048 ingress, 256 egress |
| | BlackDiamond 8900 series 8900-10G24X-c, group of 12 ports 8900 xl-series, 8900-G96T-c 8900-40G6X-xm | 1,024 ingress, 256 egress 4,096 ingress, 512 egress 512 ingress, 512 egress |
| | BlackDiamond X8 a-series modules | 512 ingress, 512 egress |
| | BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules | 8,192 ingress, 1,024 egress |
| | BlackDiamond BDXB-40G12X-XL per group of 3 ports | 8,192 ingress, 1,024 egress |
| | E4G-200 | 1,024 ingress 256 egress |
| | Summit X440, X430 per group of 24 ports | 512 ingress 2,048 ingress, |
| | Summit X460, E4G-400, per group of 24 ports | 256 egress 4,096 ingress, |
| | Summit X480 | 512 egress 512 ingress, 512 egress |
| | Summit X670 with VIM4-40G4x | 512 ingress, 512 egress |
| Summit X480 with VIM3-40G4X | 512 ingress, 512 egress | |
| Summit X460-G2, X450-G2 | 1,024 ingress, 512 egress | |
| Summit X770, X670-G2 | 1,024 ingress, 512 egress | |
| Access lists (policies) —suggested maximum number of lines in a single policy file. | All platforms | 300,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|--|
| Access lists (policies)— maximum number of rules in a single policy file. ^a | BlackDiamond 8000 series e-series, group of 24 ports c-series | 4,096 ingress, 512 egress 1,024 ingress |
| | BlackDiamond 8900 series 8900-10G24X-c, group of 12 ports 8900 xl-series 8900-40G6X-xm | 2,048 ingress, 512 egress 8,192 ingress, 1,024 egress 61,440 (up to) |
| | BlackDiamond X8 a-series modules BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules BlackDiamond BDXB-40G12X-XL per group of 3 ports E4G-200 Summit X440, X430 per group of 24 ports Summit X460 E4G-400, per group of 24 ports Summit X480 Summit X670 with VIM4-40G4x Summit X480 with VIM3-40G4X Summit X460-G2, X450-G2 Summit X770, X670-G2 | 2,048 ingress, 1,024 egress 8,192 ingress, 1,024 egress 8,192 ingress, 1,024 egress 2,048 ingress 512 egress 1,024 ingress 2,048 ingress, 256 egress 4,096 ingress, 512 egress 8,192 ingress, 1,024 egress 2,048 ingress, 1,024 egress 2,048 ingress, 1,024 egress 4,096 ingress, 1,024 egress 4,096 ingress, 1,024 egress |
| | Summit X450-G2, X460-G2, X460, X480, E4G-400 | 2,048 ingress only |
| | Summit X670-G2, X770, E4G200, X670 | 1,024 ingress only |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|--|
| Access lists (slices)—number of ACL slices. | BlackDiamond 8000 series c-series, group of 48 ports | 16 |
| | BlackDiamond 8900 series 8900 xl-series 8900-10G24X-c modules, group of 12 ports 8900-G96T-c modules, group of 48 ports 8900-40G6X-xm | 17 ^b 12 ingress, 4 egress 16 ingress, 4 egress 10 ingress, 4 egress |
| | BlackDiamond X8 a-series modules BlackDiamond X8-100G4X modules BlackDiamond XB-100G4X-XL modules E4G-200 Summit X440, X430 Summit X460, E4G-400, X460-G2, X450-G2 Summit X480 Summit X670 VIM4-40G4x Summit X480 VIM3-40G4X Summit X770, X670-G2 | 10 ingress, 4 egress 16 ingress, 4 egress 17 ingress, 4 egress 8 ingress, 4 egress 4 ingress 16 ingress, 4 egress 17 ingress ^b , 4 egress 10 ingress, 4 egress 10 ingress, 4 egress 12 ingress, 4 egress |
| Access lists (slices)—number of ACL slices in first stage (VFP). | Summit X450-G2, X460-G2, X670-G2, X770, E4G-200, E4G-400, X460, X480, X670 | 4 ingress only |
| ACL Per Port Meters—number of meters supported per port. | E4G-200 E4G-400 BlackDiamond X8, BlackDiamond 8800 Summit X430, X440 Summit X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 | 8 16 16 8 16 |
| Meters Packets-Per-Second Capable | BlackDiamond X8, BlackDiamond 8800 (8900-40G6X-c only) E4G-200, E4G-400 Summit X480 Summit X430, X440, X450-G2, X460, X460-G2, X670, X670-G2, X770 | Yes Yes No Yes |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|--|
| AVB (audio video bridging) —maximum number of active streams. Note: * It is recommended that you do not use on more than 8 ports on this switch. | Summit X440, X460, X460-G2, X450-G2 Summit X670, X670-G2 Summit X430 | 1,024 4,096 100* |
| BFD sessions —maximum number of BFD sessions. | All platforms (default timers—1 sec) BlackDiamond X8 and 8800 (minimal timers—50 msec) All Summits (minimal timers—100 msec) | 512 10 ^c 10 ^c |
| BGP (aggregates) —maximum number of BGP aggregates. | All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher | 256 |
| BGP (networks) —maximum number of BGP networks. | All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher BlackDiamond X8 series | 1024 |
| BGP (peers) —maximum number of BGP peers. Note: *With default keepalive and hold timers. | BlackDiamond X8 series, xl-series, 8000 series All Summits, except X450-G2, X480, X440, X430, E4G-200 E4G-400H Summit X480 | 512 128* 128* 512 |
| BGP (peer groups) —maximum number of BGP peer groups. | BlackDiamond 8900 series BlackDiamond 8800 BlackDiamond X8 series Summit X480 Summit X770, X670-G2, X670v-48t, X670, X460-G2, X460 (with Core license or higher) | 128 64 128 128 64 |
| BGP (policy entries) —maximum number of BGP policy entries per route policy. | All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher | 256 |
| BGP (policy statements) —maximum number of BGP policy statements per route policy. | All platforms (except E4G-200, X430, X440, and X450-G2) with Core license or higher | 1,024 |
| BGP multicast address-family routes —maximum number of multicast address-family routes. | BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8-xl series Summit X480 E4G-400 Summit X460, X460-G2, X670, X670-G2, X770 | 524,256 (up to) ^b 1,048,544 (up to) ⁱ 524,256 (up to) ^b 25,000 25,000 |
| BGP (unicast address-family routes) —maximum number of unicast address-family routes. | BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8-xl series Summit X480 E4G-400 Summit X460, X460-G2, X670, X670-G2, X770 | 524,256 (up to) ^b 1,048,544 (up to) ⁱ 524,256 (up to) ^b 25,000 25,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|--|
| BGP (non-unique routes) —maximum number of non-unique BGP routes. | BlackDiamond 8900 xl-series, X8 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series Summit X460, X460-G2, X670, X670-G2, X770 Summit X480 E4G-400 | 1,200,000 24,000 2,000,000 25,000 1,000,000 25,000 |
| BGP ECMP —maximum number of equalcost multipath for BGP and BGPv6. | All platforms, except Summit X430, X440, and E4G-200 BlackDiamond 8800 G48Te2 (for BGPv6) | 2, 4, or 8 N/A |
| BGPv6 (unicast address-family routes) —maximum number of unicast address family routes. | BlackDiamond 8900 xl-series, BlackDiamond X8 xl-series BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond X8 series Summit X460, X460-G2 Summit X480 Summit X670, X670-G2, X770 E4G-400 | 20,000 6,000 240 8,000 6,000 20,000 8,000 6,000 |
| BGPv6 (non-unique routes) —maximum number of non-unique BGP routes. | BlackDiamond 8900 xl-series BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond X8 series, X8 xl-series Summit X460, X460-G2 Summit X670, X670-G2, X770 E4G-400 | 24,000 18,000 720 24,000 18,000 24,000 18,000 |
| BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per virtual router. | All platforms, except Summit X430 | 4 |
| BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per VLAN. | All platforms, except Summit X430 | 4 |
| CES TDM pseudowires —maximum number of CES TDM pseudowires per switch. | E4G-200 and E4G-400 | 256 |
| Connectivity fault management (CFM) —maximum number of CFM domains. Note: With Advanced Edge license or higher. | All platforms | 8 |
| CFM —maximum number of CFM associations. Note: With Advanced Edge license or higher. | All platforms | 256 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|---|
| <p>CFM—maximum number of CFM up end points.</p> <p>Note: With Advanced Edge license or higher.</p> | BlackDiamond 8000 series, X8 series Summit series | 32 |
| <p>CFM—maximum number of CFM down end points.</p> <p>Note: With Advanced Edge license or higher.</p> | BlackDiamond 8000 series, X8 series Summit series X460, E4G-200, E4G-400 (non-load shared ports) All other platforms | 32 256 (non-load shared ports), 32 (load shared ports) 32 |
| <p>CFM—maximum number of CFM remote end points per up/down end point.</p> <p>Note: With Advanced Edge license or higher.</p> | All platforms | 2,000 |
| <p>CFM—maximum number of dot1ag ports.</p> <p>Note: With Advanced Edge license or higher.</p> | All Summits, except X430, X450-G2 | 128 |
| <p>CFM—maximum number of CFM segments.</p> <p>Note: With Advanced Edge license or higher.</p> | All platforms | 1,000 |
| <p>CFM—maximum number of MIPs.</p> <p>Note: With Advanced Edge license or higher.</p> | All platforms | 256 |
| <p>CLEAR-Flow—total number of rules supported. The ACL rules plus CLEAR-Flow rules must be less than the total number of supported ACLs.</p> | BlackDiamond X8, BlackDiamond 8800 Summit X440, X430 Summit X670 Summit X460, X460-G2, X770, X670-G2, X450-G2 Summit X480 E4G-200 E4G-400 | 4,096 1,024 2,048 4,094 8,192 2,048 4,094 |
| <p>Data Center Bridging eXchange (DCBX) protocol Type Length Value (TLVs)—maximum number of DCBX application TLVs.</p> | All platforms | 8 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|---|
| 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) |
| DHCP snooping entries —maximum number of DHCP snooping entries. | All Summits BlackDiamond X8 | 2,048 6,000 |
| Dynamic ACLs —maximum number of ACLs processed per second. Note: Limits are load dependent. | Summit X480, X670 with 50 DACLs with 500 DACLs | 10 5 |
| | BlackDiamond X8 BlackDiamond 8800 | N/A N/A |
| EAPS domains —maximum number of EAPS domains. Note: An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains. | BlackDiamond 8000 series, X8 series Summit X670-G2, X450-G2, and X770 Summit X670, X480, X460, X460-G2, X440, E4G-200, E4G-400 Summit X430 | 64 64 32 4 |
| EAPsv1 protected VLANs —maximum number of protected VLANs. | BlackDiamond 8000 series, X8 series Summit series, E4G-200, E4G-400 | 2,000 1,000 |
| EAPsv2 protected VLANs —maximum number of protected VLANs. | BlackDiamond 8000 series, X8 series All Summits (except X430, X440), E4G-200, E4G-400 | 2,000 500 |
| ELSM (vlan-ports) —maximum number of VLAN ports. | BlackDiamond 8000 series, X8 series All Summits, E4G-200, E4G-400 | 5,000 |
| ERPS domains —maximum number of ERPS domains without CFM configured. | BlackDiamond 8800 series, X8 series Summit series (except X430), E4G-200, E4G-400 Summit X430 | 32 32 4 |
| ERPS domains —maximum number of ERPS domains with CFM configured. | BlackDiamond 8800 series, X8 series E4G-200, E4G-400 Summit X460 Summit X430 Summit X440, X770, X670, X670-G2, X480, X460-G2, X450-G2 | 16 32 32 4 16 |
| ERPSv1 protected VLANs —maximum number of protected VLANs. | BlackDiamond 8800 series, X8 series All Summits, E4G-200, E4G-400 | 2,000 1,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|--|
| ERPSv2 protected VLANs —maximum number of protected VLANs. | BlackDiamond 8800 series, X8 series All Summits (except X430), E4G-200, E4G-400 | 2,000 500 |
| ESRP groups —maximum number of ESRP groups. | All platforms | 7 |
| ESRP domains —maximum number of ESRP domains. | All platforms | 64 |
| ESRP VLANs —maximum number of ESRP VLANs. | BlackDiamond 8800 BlackDiamond X8 All Summits E4G-200. E4G-400 | 1,000 2,048 1,000 1,000 |
| ESRP (maximum ping tracks) —maximum number of ping tracks per VLAN. | All platforms (except Summit X430) | 8 |
| ESRP (IP route tracks) —maximum IP route tracks per VLAN. | All platforms (except Summit X430) | 8 |
| ESRP (VLAN tracks) —maximum number of VLAN tracks per VLAN. | All platforms (except Summit X430) | 1 |
| Forwarding rate —maximum L3 software forwarding rate. | BlackDiamond 8000 series BlackDiamond X8 series Summit X770 Summit X670-G2 Summit X670 Summit X480 Summit X460-G2 Summit X460 Summit X450-G2 Summit X440 E4G-200 E4G-400 | 10,000 pps 20,000 pps 14,000 pps 25,000 pps 14,829 pps 14,509 pps 28,000 pps 5,222 pps 27,000 pps 5,418 pps 8,718 pps 5,536 pps |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|--|
| FDB (unicast blackhole entries)—maximum number of unicast blackhole FDB entries. | BlackDiamond 8900 series 8900 c-series 8900 xl-series 8900-40G6X-xm | 32,000 524,288 (up to) ^b 128,000 |
| | BlackDiamond 8000 e-series BlackDiamond 8800 c-series BlackDiamond X8 a-series modules BlackDiamond X8-100G4X modules. BlackDiamond X8 xl-series module E4G-200, E4G-400 Summit X440, X430 Summit X480 Summit X460 Summit X460-G2 Summit X670 VIM4-40G4x, X480 VIM3-40G4X Summit X770, X670-G2 Summit X670, X670v-48t Summit X450-G2 | 8,000 32,000 128,000 384,000 384,000 ^d 32,000 16,000 524,288 (up to) ^b 32,000 49,152 ^e 128,000 294,912 130,000 ^e 34,000 |
| FDB (multicast blackhole entries)—maximum number of multicast blackhole FDB entries. | BlackDiamond 8000 series, X8 series Summit X480, X460-G2, X460, X440, X430, X450-G2 Summit X770, X670, X670-G2, X670v-48t, X480 VIM3-40G4X E4G-200, E4G-400 | 1,024 1,024 4,096 1,024 |
| FDB (maximum L2 entries)—maximum number of MAC addresses. | BlackDiamond 8000 c-series BlackDiamond 8000 e-series BlackDiamond 8000 (system), except 8900 xl-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series modules BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X440, X430 Summit X480 (40G4X) Summit X460 Summit X670-G2 Summit X460-G2 Summit X670 Summit X770 Summit X450-G2 | 32,768 ^f 8,192 ^f 128,000 ^f 524,488 (up to) ^b 128,000 ^f 384,000 ^f 1,048,576 (up to) ^{bg} 32,000 ^f 16,000 ^f 524,488 (up to) ^{bf} 32,000 ^f 294,912 ^f 96,000 ^f 128,000 ^{ef} 294,912 ^f 68,000 |
| FDB (Maximum L2 entries)—maximum number of multicast FDB entries. | BlackDiamond X8, 8800 Summit X770, X670, X670-G2 Summit X480, X460, X460-G2, X430, X440, X450-G2 E4G-200, E4G-400 | 1,024 4,096 1,024 1,024 |
| FIP Snooping VLANs | BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) Summit X670 | 768 |
| FIP Snooping Virtual Links (FPMA mode) per port group | BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) Summit X670 | 1,908 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|------------|
| FIP Snooping FCFs (with perimeter port) per port group | BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) | 238 |
| FIP Snooping FCFs (with Enode-to-FCF port) | BlackDiamond X8 BlackDiamond 8800 (8900-40G6X-c only) Summit X670 | 212 |
| Identity management— maximum number of Blacklist entries. | All platforms, except Summit X430 Summit X430 | 512 N/A |
| Identity management— maximum number of Whitelist entries. | All platforms, except Summit X430 Summit X430 | 512 N/A |
| Identity management— maximum number of roles that can be created. | All platforms, except Summit X430 Summit X430 | 64 N/A |
| Identity management— maximum role hierarchy depth allowed. | All platforms, except Summit X430 Summit X430 | 5 N/A |
| Identity management— maximum number of attribute value pairs in a role match criteria. | All platforms, except Summit X430 Summit X430 | 16 N/A |
| Identity management— maximum of child roles for a role. | All platforms, except Summit X430 Summit X430 | 8 N/A |
| Identity management— maximum number of policies/ dynamic ACLs that can be configured per role. | All platforms, except Summit X430 Summit X430 | 8 N/A |
| Identity management— maximum number of LDAP servers that can be configured. | All platforms, except Summit X430 Summit X430 | 8 N/A |
| Identity management— maximum number of Kerberos servers that can be configured. | All platforms, except Summit X430 Summit X430 | 20 N/A |
| Identity management— maximum database memory-size. | All platforms, except Summit X430 Summit X430 | 512 N/A |
| Identity management— recommended number of identities per switch. | All platforms, except Summit X430 Summit X430 | 100 N/A |
| Note: Number of identities per switch is for a default identity management database size (512 Kbytes) across all platforms. | | |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|--|
| Identity management —recommended number of ACL entries per identity. Note: Number of ACLs per identity based on system ACL limitation. | All platforms, except Summit X430 Summit X430 | 20 N/A |
| 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 Summit X430 Summit X430 | 500 N/A |
| IGMP sender —maximum number of IGMP senders per switch (IP multicast compression enabled). ^m Note: Assumes source-group-vlan mode. For additional limits, see: <ul style="list-style-type: none"> Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900-10G24X-c modules BlackDiamond 8900-G96T-c modules BlackDiamond 8900-40G6X-xm BlackDiamond 8900 xl-series BlackDiamond X8 a-series modules BlackDiamond X8 b-series modules E4G-200 E4G-400 Summit X440 Summit X460 Summit X460-G2 Summit X450-G2 Summit X480 Summit X670 Summit X770, X670-G2 Summit X430 | 2,048 ⁱ 500 ⁱ 2,048 ⁱ 4,096 ⁱ 3,000 ⁱ 12,000 ⁱ 4,096 ^b 64,000 ⁱ 3,000 ^{i j} 6,000 ^{i j} 192 ⁱ 6,000 ⁱ 30,000 ^h 21,000 ^h 12,000 ⁱ 3,000 77,500 ⁱ 192 |
| IGMP snooping per VLAN filters —maximum number of VLANs supported in per-VLAN IGMP snooping mode. | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series modules BlackDiamond X8 b-series modules E4G-200, E4G-400 Summit X440 Summit X460, X670, X440 Summit X460-G2 Summit X450-G2 Summit X480 Summit X770, X670-G2 | 2,000 448 1,000 4,000 1,000 1,000 4,000 1,000 448 1,000 1,500 2,048 4,000 2,000 |
| IGMPv1/v2 SSM-map entries —maximum number of IGMPv1/v2 SSM mapping entries. | All platforms | 500 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|---|
| IGMPv1/v2 SSM-MAP entries —maximum number of sources per group in IGMPv1/v2 SSM mapping entries. | All platforms | 50 |
| IGMPv2 subscriber —maximum number of IGMPv2 subscribers per port. ⁿ | BlackDiamond 8800 c-series, 8900 c-series, X8 series Summit X430, X460, E4G-200, E4G-400, X440 Summit X480, X670, X670v-48t Summit X770, X670-G2, X460-G2, X450-G2 | 2,000 1,000 2,000 4,000 |
| IGMPv2 subscriber —maximum number of IGMPv2 subscribers per switch. ⁿ | BlackDiamond 8800 c-series, 8900 c-series, X8 series Summit X430, X440, E4G-200 Summit X770, X670-G2 Summit X460, X460-G2, X480, X670, E4G-400, X670v-48t, X450-G2 | 20,000 10,000 30,000 20,000 |
| 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 | BlackDiamond 8800 e-series BlackDiamond 8800 c-series BlackDiamond 8900 series BlackDiamond X8 series Summit X480, X670, X670v-48t, E4G-200, X440 Summit X770, X670-G2, X460-G2, X450-G2 Summit X460, E4G-400 | 1,000 2,000 5,000 3,000 1,000 4,000 2,000 |
| IGMPv3 subscriber —maximum number of IGMPv3 subscribers per switch. ⁿ | BlackDiamond 8800 e-series BlackDiamond 8800 c-series BlackDiamond 8900 series BlackDiamond X8 series Summit X670, X670v-48t, X480, E4G-200, X440 Summit X460, X460-G2, E4G-400, X450-G2 Summit X770, X670-G2 | 10,000 20,000 30,000 20,000 10,000 20,000 30,000 |
| IP ARP entries in software —maximum number of IP ARP entries in software. Note: May be limited by hardware capacity of FDB (maximum L2 entries). | BlackDiamond X8-100G4X modules Summit X670-G2, X770 Summit X670, X480, X460, X440, X430 Summit X460-G2 Summit X450-G2 E4G-200, E4G-400 | 229,374 (up to) ^h 131,072 (up to) ^h 20,480 ^h 57,344 (up to) ^h 47,000 (up to) ^h 20,480 |
| IP ARP entries in software with distributed mode on —maximum number of IP ARP entries in software with distributed mode on. | BlackDiamond 8000 series with 8900-MSM128, MSM-48c, or MSM-96 and only 8900 xl-series I/O modules BlackDiamond 8000 series with any I/O modules that are not 8900 xl-series BlackDiamond X8 series All other platforms | 260,000 100,000 172,000 N/A |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|--------------------------------|
| IPv4 ARP entries in hardware with distributed mode on —maximum number of IP ARP entries in hardware with distributed mode on | Per BlackDiamond 8900-10G8X-xl, up to 260,000 per system | 32,500 ^b |
| | Per BlackDiamond 8900-G48X-xl or 8900-G48T-xl, up to 130,000 per system | 16,250 ^b |
| | Per BlackDiamond 8000 c-series, up to 18,000 per system | 8,000 |
| | Per BlackDiamond 8900-40G6X-xm, up to 22,000 per system | 8,000 |
| | Per BlackDiamond X8 a-series, up to 28,000 per system | 12,000 |
| | Per BlackDiamond X8 xl-series, up to 172,000 per system | 172,000 |
| | All other platforms | N/A |
| IPv4 ARP entries in hardware with minimum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with minimum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series switches, assumes number of IP route reserved entries is 100 or less. | BlackDiamond 8800 c-, xm-series | 8,000 |
| | BlackDiamond 8000 e-series | 1,000 ⁱ |
| | BlackDiamond 8900 xl-series | 16,000 |
| | BlackDiamond X8 a-series | 16,000 |
| | BlackDiamond X8-100G4X modules | 182,000 (up to) ^{h,m} |
| | BlackDiamond X8 xl-series | 294,000 (up to) ⁱ |
| | E4G-200 | 8,000 |
| | E4G-400 | 16,000 |
| | Summit X440 | 412 |
| | Summit X670, X480 (40G4X) | 8,000 |
| | Summit X460, X480 | 16,000 |
| Summit X460-G2 | 50,000 (up to) ^h | |
| Summit X770, X670-G2 | 108,000 (up to) ^h | |
| Summit X450-G2 | 39,000 (up to) ^h | |
| IPv4 ARP entries in hardware with maximum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with maximum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is “maximum.” | BlackDiamond 8800 c-, xm-series | 6,000 ⁱ |
| | BlackDiamond 8000 e-series | 500 ⁱ |
| | BlackDiamond 8900 xl-series | 12,000 ⁱ |
| | BlackDiamond X8 a-series | 12,000 ⁱ |
| | BlackDiamond X8-100G4X modules | 172,000 (up to) ^{h,j} |
| | BlackDiamond X8 xl-series | 290,000 (up to) ⁱ |
| | E4G-200 | 6,000 ⁱ |
| | E4G-400 | 12,000 ⁱ |
| | Summit X440 | 380 |
| | Summit X460, X480 | 12,000 ⁱ |
| | Summit X670, X480 VIM3-40G4X | 6,000 ⁱ |
| Summit X770, X670-G2 | 98,000 (up to) ^h | |
| Summit X460-G2 | 43,000 (up to) ^h | |
| Summit X450-G2 | 29,000 (up to) ^h | |
| IP flow information export (IPFIX) —number of simultaneous flows. | BlackDiamond 8900 xl-series modules | 4,096 ingress, 4,096 egress |
| | BlackDiamond 8900 c-series modules | 4,096 ingress, 4,096 egress |
| | BlackDiamond X8 b-series modules | 2,048 ingress, 2,048 egress |
| | Summit X460-24t/x/p, X460-G2 | 2,048 ingress, 2,048 egress |
| | Summit X480, X460-48t/x/p | 4,096 ingress, 4,096 egress |
| | E4G-400 | 2,048 ingress, 2,048 egress |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|---------------------------------|
| 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. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is 0, and number of IPv4 ARP entries present is 100 or less. | BlackDiamond 8800 c-series | 18,000 ⁱ |
| | BlackDiamond 8000 e-series | 1,000 ⁱ |
| | BlackDiamond 8900 xl-series | 40,000 ^b |
| | BlackDiamond 8900-40G6X-xm | 22,000 ⁱ |
| | BlackDiamond X8 a-series | 28,000 ⁱ |
| | BlackDiamond X8-100G4X and X8 xl-series | 311,000 (up to) ^h |
| | E4G-200 | 18,000 ⁱ |
| | E4G-400 | 20,000 ⁱ |
| | Summit X440 | 448 |
| | Summit X460 | 20,000 ⁱ |
| | Summit X460-G2 | 73,000 ^h |
| | Summit X480 | 40,000 ^b |
| Summit X670, X480 VIM3-40G4X | 22,000 ⁱ | |
| Summit X770, X670-G2 | 176,000 (up to) ^h | |
| Summit X450-G2 | 61,000 (up to) ^h | |
| IPv4 routes —maximum number of IPv4 routes in software (combination of unicast and multicast routes). | BlackDiamond 8900 xl-series with 8900-MSM128, MSM-48c, or MSM-96 | 524,256 (up to) ^b |
| | All other BlackDiamond 8000 series hardware | 25,000 |
| | BlackDiamond X8 series | 1,048,544 (up to) ⁱ |
| | BlackDiamond X8 with BDX X8 xl-series | 256 |
| | Summit X440 | 25,000 |
| | Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 | 524,256 (up to) ^b |
| | Summit X480 | 25,000 |
| IPv4 routes (LPM entries in hardware) — number of IPv4 routes in hardware. | BlackDiamond 8800 c-series | 12,000 |
| | BlackDiamond 8000 e-series | 480 |
| | BlackDiamond 8900 xl-series | 524,256 (up to) ^b |
| | BlackDiamond 8900-40G6X-xm | 16,000 ^e |
| | BlackDiamond X8 series | 16,000 ^e |
| | BlackDiamond BDX X8 xl-series | 1,048,544 (up to) |
| | E4G-200, E4G-400 | 12,000 |
| | Summit X440 | 32 |
| | Summit X460, X460-G2 | 12,000 |
| | Summit X480 | 524,256 (up to) ^{b, o} |
| | Summit X480 VIM3-40G4X | 16,000 ^o |
| | Summit X670 | 12,000 |
| Summit X770, X670-G2, X450-G2 | 16,000 | |
| 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. | BlackDiamond 8000 series | 512 |
| | BlackDiamond X8 series | 2,048 |
| | E4G-200, E4G-400 | 512 |
| | Summit X440 | 254 |
| | Summit X460, X480 | 512 |
| | Summit X770, X670, X670-G2, X460-G2, X450-G2 | 2,048 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|--|
| IPv6 host entries in hardware —maximum number of IPv6 neighbor entries in hardware. | BlackDiamond 8800 c-, xm-series BlackDiamond 8000 e-series BlackDiamond 8900-10G24X-c modules BlackDiamond 8900-G96T-c modules BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X BlackDiamond X8 xl-series E4G-200 E4G-400 Summit X440 Summit X460, X670, X480 VIM3-40G4X Summit X770, X670-G2 Summit X480, X670v-48t Summit X460-G2 Summit X450-G2 | 3,000 ⁱ 250 ⁱ 2,000 ⁱ 4,000 ⁱ 8,192 (up to) ^{bi} 3,000 ⁱ 49,000 ^{ih} 49,000 ^{il} 2,000 ⁱ 3,000 ⁱ 192 ^l 3,000 ⁱ 36,750 ⁱ 6,000 ⁱ 22,000 ⁱ 12,000 ⁱ |
| IPv6 routes (LPM entries in hardware) —maximum number of IPv6 routes in hardware. | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 xm-series BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X440 Summit X460, X460-G2 Summit X480 Summit X670, X480 (VIM3-40G4X), X670-G2, X770, X450-G2 | 6,000 240 8,000 245,760 (up to) ^b 8,000 524,288 (up to) ^l 6,000 16 6,000 245,760 (up to) ^b 8,000 |
| IPv6 routes with a mask greater than 64 bits in hardware —maximum number of such IPv6 LPM routes in hardware. | BlackDiamond 8000 c-, e-, xm-series BlackDiamond 8000 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series E4G-200, E4G-400 Summit X480 Summit X440, X460, X460-G2, X670, X670-G2, X770, X480 (VIM3-40G4X), X450-G2 | 256 245,760 (up to) ^k 256 524,288 (up to) ^l 256 245,760 (up to) ^k 256 |
| IPv6 route sharing in hardware —route mask lengths for which ECMP is supported in hardware. Note: * >64 single path only | Summit X460, X480, X670, X670V-48t, X450-G2 E4G-200, E4G-400 BlackDiamond 8800 (all I/O modules, except G48Te2) Summit X460-G2, X670-G2, X770 BlackDiamond X8 a-series BlackDiamond X8-100G4X modules BlackDiamond X8 xl-series Summit X440, X430 BlackDiamond 8800 G48Te2 | 0-128 0-128 0-128 0-64 * 0-128 0-64 * 0-128 ^l N/A N/A |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|--|
| IPv6 routes in software—maximum number of IPv6 routes in software. | BlackDiamond 8900 xl-series with 8900-MSM128, MSM-48c, or MSM-96 | 245,760 (up to) ^k 25,000 |
| | All other BlackDiamond 8000 series hardware | 25,000 |
| | BlackDiamond X8 series | 524,288 (up to) ^l |
| | BlackDiamond X8 with xl-series | 25,000 |
| | Summit X460, X460-G2, X670, X670-G2, X770, X450-G2, E4G-200, E4G-400 | 245,760 (up to) ^k 256 |
| IP router interfaces—maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs. | Summit X460-G2, X670, X770, X670-G2, X450-G2, and BlackDiamond X8 | 2,048 |
| | BlackDiamond 8800 | 512 |
| | Summit X440 | 254 |
| | Summit X480, X460 E4G-200, E4G-400 | 512 512 |
| IP multicast static routes—maximum number of permanent multicast IP routes. | All platforms (except Summit X430, X440) | 1,024 |
| | Summit X430, X440 | 32 |
| IP unicast static routes—maximum number of permanent IP unicast routes. | All platforms (except Summit X430, X440) | 1,024 |
| | Summit X430, X440 | 32 |
| IP route sharing (maximum gateways)—Configurable maximum number of gateways used by equal cost multipath OSPF, BGP, IS-IS, static routes, or L2VPNs. Routing protocol OSPF is limited to 16 ECMP gateways per destination. Routing protocols BGP and IS-IS are limited to 8 ECMP gateways per destination. Static routes are limited to 32 next-hops. 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. | All platforms, except Summit X430, X440, X670, and BlackDiamond X8 | 2, 4, 8, 16, 32 |
| | Summit X670, BlackDiamond X8 | 2, 4, 8, 16, or 32, |
| | Summit X430, X440 | or 64 |
| | BlackDiamond 8800 G48Te2 (for IPv6) | N/A N/A |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|--------------------------------|
| IP route sharing (total destinations)—maximum number of unique destinations used by multipath OSPF, OSPFv3, BGP, IS-IS, or static routes. | BlackDiamond 8800 c-series | 12,256 |
| | BlackDiamond 8000 e-series | 480 |
| | BlackDiamond 8900 xl-series | 524,256 (up to) ^b |
| | BlackDiamond 8900-40G6X-xm | 16,352 |
| | BlackDiamond X8 | 16,352 |
| | BlackDiamond X8 xl-series | 1,048,544 (up to) ⁱ |
| | E4G-200, E4G-400 | 12,256 |
| | Summit X480 | 524,256 (up to) ^b |
| | Summit X670, X670-G2, X770, X450-G2, X480 (VIM3-40G4X) | 16,352 |
| | Summit X460-G2, X460 | 12,256 |
| | <p>Note:</p> <p>For platforms with limit of 524,256 or higher, the total number of "destination+gateway" pairs is limited to 2,097,024. For example, if the number of unique destinations is 524,256, only 2 gateways per destination is supported.</p> <p>For other platforms, each limit is based on up to 8 gateways per destination for BGP and IS-IS routing protocols, up to 16 gateways per destination for OSPF, or up to 32 gateways per destination for static routes.</p> | |

Table 4: Supported Limits (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. | BlackDiamond 8800 c-, xl-, and xm-series default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 | 510 1,022 254 126 62 |
| | BlackDiamond 8000 e-series default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 | 30 62 14 6 2 |
| | BlackDiamond X8 series, Summit X670 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 if maximum gateway is 64 | 510 1,022 254 126 62 30 |
| | Summit X460, X460-G2, X450-G2, X480, X670, X670-G2, X770, E4G-200, E4G-400 default maximum gateways of 4 if maximum gateways is 2 if maximum gateways is 8 if maximum gateways is 16 if maximum gateways is 32 | 510 1,022 254 126 62 |
| IP multinetting (secondary IP addresses) —maximum number of secondary IP addresses per VLAN. | BlackDiamond 8800, BlackDiamond X8 All Summits, except X440, X430 Summit X440 | 64 255 32 |
| IS-IS adjacencies —maximum number of supported IS-IS adjacencies. | BlackDiamond 8000 series, BlackDiamond X8 series BlackDiamond 8900 xl-series Summit X440, X460, X460-G2, X480, X670, X670-G2, X770 Summit X450-G2 E4G-200 E4G-400 | 128 255 128 N/A 256 128 |
| IS-IS ECMP —maximum number of equal cost multipath for IS-IS. | All platforms, except Summit X440, X430 BlackDiamond 8800 G48Te2 (for IPv6) | 2, 4, or 8 N/A |
| IS-IS interfaces —maximum number of interfaces that can support IS-IS. | All platforms, except Summit X440, x430 | 255 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|------------------|
| IS-IS routers in an area—recommended maximum number of IS-IS routers in an area. | Summit X480 | 128 |
| | All other platforms, except Summit X440, X430 | 256 |
| IS-IS route origination—recommended maximum number of routes that can be originated by an IS-IS node. | BlackDiamond 8000 series, BlackDiamond X8 series | 20,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 30,000 |
| | E4G-400 | 25,000 |
| | E4G-200 | 20,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770, X480 | 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. | BlackDiamond 8000 series, BlackDiamond X8 series | 25,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 120,000 |
| | Summit X480 | 50,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 | 25,000 |
| | E4G-200, E4G-400 | 25,000 |
| IS-IS IPv4 L2 routes—recommended maximum number of IS-IS Level 2 routes. | BlackDiamond 8000 series | 20,000 |
| | BlackDiamond X8 series | 25,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 120,000 |
| | Summit X480 | 50,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 E4G-200, E4G-400 | 25,000 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. | BlackDiamond 8000 series, X8 series, 8900 xl-series | 20,000 |
| | E4G-200, E4G-400 | |
| | Summit X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 | |
| 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. | BlackDiamond 8000 series, X8 series | 10,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 40,000 |
| | Summit X480 | 25,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770, E4G-400 | 10,000 |
| | | |
| IS-IS IPv6 L2 routes—recommended maximum number of IS-IS Level 2 routes. | BlackDiamond 8000 series, X8 series | 10,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 40,000 |
| | Summit X480 | 15,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 | 10,000 |
| | E4G-200, E4G-400 | 10,000 |
| IS-IS IPv6 L1 routes in an L1/L2 router—recommended maximum number of IS-IS Level 1 routes in a L1/L2 router. | BlackDiamond 8000 series, X8 series | 10,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 15,000 |
| | Summit X480 | 15,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770, E4G-400 | 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. | BlackDiamond 8000 series, X8 series | 20,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 60,000 |
| | Summit X480 | 40,000 |
| | Summit X450-G2, X460, X460-G2, X670, X670-G2, X770 | 20,000 |
| | E4G-200, E4G-400 | 20,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|------------------------|
| 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. | BlackDiamond 8000 series, X8 series | 20,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 60,000 |
| | Summit X480 | 40,000 |
| | Summit X450-G2, X460,X460-G2, X670, X670-G2, X770 | 20,000 |
| | E4G-200, E4G-400 | 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. | BlackDiamond 8000 series | 20,000 |
| | BlackDiamond X8 series | |
| | BlackDiamond 8900 xl-series | |
| | Summit X450-G2, X460, X460-G2, X480, X670, X670-G2, X770 | |
| | E4G-200, E4G-400 | |
| Jumbo frames —maximum size supported for jumbo frames, including the CRC. | All platforms | 9,216 |
| L2 VPN: VCCV (pseudowire Virtual Circuit Connectivity Verification) VPNs per switch —maximum number of VCCV enabled VPLS VPNs. | All platforms, except Summit X440, X430, and X450-G2 | 16 |
| L2 VPN: VPLS MAC addresses —maximum number of MAC addresses learned by a switch. | BlackDiamond 8900 xl-series | 512,000 |
| | BlackDiamond 8900-40G6X-xm | 128,000 |
| | BlackDiamond X8 a-series modules | 128,000 |
| | BlackDiamond X8-100G4X modules | 384,000 |
| | BlackDiamond X8 xl-series | 1,048,576 ⁹ |
| | E4G-200, E4G-400 | 32,000 |
| | Summit X460 | 32,000 |
| | Summit X480 | 512,000 |
| | Summit X670, Summit X670V-48t, Summit X770 | 128,000 |
| | Summit X480 (40G VIM) | 121,000 |
| | Summit X670-G2 | 140,000 |
| Summit X460-G2 | 55,000 | |
| L2 VPN: VPLS VPNs —maximum number of VPLS virtual private networks per switch. | BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series E4G-200, E4G-400 Summit X460, X460-G2, X480, X670, X670V-48t, X480 (40G VIM), X770, X670-G2 | 1,023 |
| L2 VPN: VPLS peers —maximum number of VPLS peers per VPLS instance. | BlackDiamond 8900 xl-series, 8900-40G6x-xm, X8 series | 64 |
| | Summit X770, X670-G2, X670v-48t, X480, X460-G2 | 64 |
| | Summit X670, X460 | 32 |
| | E4G-200, E4G-400 | 32 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|---------------------|
| L2 VPN: LDP pseudowires — maximum number of pseudowires per switch. | BlackDiamond 8900 xl-series, X8 series | 7,000 |
| | BlackDiamond 8900-40G6X-xm | 3,000 |
| | E4G-200, E4G-400 | 1,000 |
| | Summit X770 | 7,800 |
| | Summit X670-G2, X670v-48t, X480 | 7,000 |
| | Summit X670 | 3,000 |
| | Summit X460-G2 Summit X460 | 7,116 1,000 |
| L2 VPN: static pseudowires — maximum number of static pseudowires per switch. | BlackDiamond 8900 xl-series, X8 series | 7,116 |
| | BlackDiamond 8900-40G6X-xm | 3,020 |
| | Summit X460, X480, X670V-48t | 7,116 |
| | Summit X770 | 15,308 |
| | Summit X480-40G, Summit X670 | 3,020 |
| | Summit X670-G2, X460-G2 | 7,000 |
| | E4G-200 E4G-400 | 2,764 6,860 |
| L2 VPN: Virtual Private Wire Service (VPWS) VPNs — maximum number of virtual private networks per switch. | Summit X460 | 1,000 |
| | Summit X480, X770 | 4,000 |
| | Summit X480-40G VIM, X670 | 2,047 |
| | Summit X670V-48t | 4,000 |
| | BlackDiamond 8900 xl-series | 4,000 |
| | BlackDiamond 8900-40G6X-xm | 2,047 |
| | BlackDiamond X8 series | 4,000 |
| | Summit X670-G2 | 4,090 |
| | Summit X460-G2 E4G-200, E4G-400 | 1,023 1,000 |
| Layer-2 IPMC forwarding caches —(IGMP/MLD/PIM snooping) in mac-vlan mode. Note: IPv6 and IPv4 L2 IPMC scaling is the same for this mode. | BlackDiamond 8800 e-series switches | 2,000 |
| | BlackDiamond 8800 c- and xl-series switches | 8,000 |
| | BlackDiamond 8800 xm-series switches | 15,000 |
| | BlackDiamond X8 series switches | 15,000 |
| | E4G-200, E4G-400 | 8,000 |
| | Summit X480, X460 | 8,000 |
| | Summit X670, X670V | 15,000 |
| | Summit X440 | 5,000 |
| | Summit X770, X670-G2 | 77,500 ^h |
| | Summit X460-G2 | 32,000 ^h |
| | Summit X430 Summit X450-G2 | 5,000 20,000 |
| Layer-2 IPMC forwarding caches —(IGMP/MLD/PIM snooping) in mixed-mode. Note: IPv6 and IPv4 L2 IPMC scaling is the same for this mode. | BlackDiamond 8800 e-series switches | N/A |
| | BlackDiamond 8800 xl- and c-series switches | 8,000 |
| | BlackDiamond 8800 xm-series switches | 15,000 |
| | BlackDiamond X8, Summit X670, X670V | 15,000 |
| | E4G-200 and E4G-400 cell site routers, Summit X460 | 8,000 |
| | Summit X440 | 5,000 |
| | Summit X770, X670-G2 | 77,500 ^h |
| | Summit X460-G2 | 24,000 |
| | Summit X480 Summit X450-G2 | 8,000 15,000,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|---------------------|
| Layer-3 IPMC forwarding caches —(PIM, MVR, PVLAN) in mixed-mode. ⁱ Note: IPv6 L3 IPMC scaling is 50% of these limits in this mode. | BlackDiamond 8800 e-series switches | N/A |
| | BlackDiamond 8800 xl- and c-series switches | 6,000 |
| | BlackDiamond 8800 xm-series switches | 3,000 |
| | BlackDiamond X8 a-series modules | 6,000 |
| | BlackDiamond X8-100G4X and modules | 64,000 |
| | E4G-200 cell site routers, Summit X670 | 3,000 |
| | E4G-400 cell site routers, Summit X460, X480, X670V | 6,000 |
| | Summit X440 | 192 |
| | Summit X770, X670-G2 | 77,500 ^h |
| Summit X450-G2 | 21,000 ^h | |
| Summit X460-G2 | 26,000 ^h | |
| 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. | BlackDiamond 8000 series without 8900-40G6X-xm | |
| | With distributed IP ARP mode off (default) | 128 |
| | With distributed IP ARP mode on | 64 |
| | BlackDiamond 8000 series 8900-40G6X-xm using address-based custom algorithm | |
| | With distributed IP ARP mode off (default) | 128 |
| | With distributed IP ARP mode on | 64 |
| | BlackDiamond 8000 series with 8900-40G6X-xm with L2, L3 or L3_L4 algorithm configured for any group | |
| | With distributed IP ARP mode off (default) | 127 |
| With distributed IP ARP mode on | 63 | |
| SummitStack with X670 with L2, L3 or L3_L4 algorithm configured for any group ^f | 127 | |
| All other SummitStack configurations and Summit series switches | 128 | |
| BlackDiamond X8 series using address-based custom algorithm | | |
| With distributed IP ARP mode off (default) | 384 | |
| With distributed IP ARP mode on | 384 | |
| BlackDiamond X8 series with L2, L3 or L3_L4 algorithm configured for any group | | |
| With distributed IP ARP mode off (default) | 127 | |
| With distributed IP ARP mode on | 63 | |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|--|
| Load sharing —maximum number of ports per load-sharing group. Note: *For custom algorithm ** For L2 and L3 algorithms Note: For a mix of Summit X770 and Summit X670 series switches in a stack, the limits are the Summit X670 limits. | BlackDiamond X8 series Summit X460-G2 (standalone) | 64 32 |
| | Summit X670 (standalone) | 32 * 16 ** |
| | Summit X670 (stacked) Summit X670-G2 (stacked) | 64 * 16 ** |
| | Summit X770 (standalone) Summit X670-G2 (standalone) Summit X460-G2 (standalone) Summit X450-G2 (standalone) | 32 |
| | Summit X770 (stacked) Summit X670-G2 (stacked) Summit X460-G2 (stacked) Summit X450-G2 (stacked) | 64 |
| | All other Summit series, SummitStacks, E4G cell site routers, and BlackDiamond 8000 series switches | 8 |
| Logged messages —maximum number of messages logged locally on the system. | All platforms | 20,000 |
| MAC address learning rate —hardware learning rat. | E4G-200 | 22 msec |
| 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 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|---|
| Maximum mirroring instances Note: The Summit X430 can only support one egress mirroring instance. | All platforms Note: Only two or four mirroring instance will be active at a time depending on the mirroring filter added to it. There are four hardware resource slots. Each single instance uses one such slot, while each ingress plus egress instance uses two slots. So this allows the you to use a total of four slots, while there are no more then two egress instances. The maximum possible combination for mirroring instances: 1 4 ingress 2 3 ingress + 1 egress 3 2 ingress + 2 egress 4 2 (ingress + egress) 5 1 (ingress + egress) + 2 ingress 6 1 (ingress + egress) + 1 egress + 1 ingress | 16 (including default mirroring instance) |
| Mirroring (filters) —maximum number of mirroring filters. Note: This is the number of filters across all the active mirroring instances. | BlackDiamond 8000 series BlackDiamond X8 series All Summit series E4G cell site routers | 128 |
| 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 | BlackDiamond 8000 series BlackDiamond X8 series All Summit series E4G cell site routers | 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. | BlackDiamond 8000 series BlackDiamond X8 series All Summit series (except X430) E4G cell site routers | 768 |
| MLAG peers —maximum number of MLAG peers allowed. | All platforms, except Summit X430 | 2 |
| MPLS RSVP-TE interfaces —maximum number of interfaces. | All platforms, except Summit X450-G2, X440, and X430 | 32 |
| MPLS RSVP-TE ingress LSPs —maximum number of ingress LSPs. | All platforms, except Summit X450-G2, X440, and X430 | 2,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|---|---|
| MPLS RSVP-TE egress LSPs— maximum number of egress LSPs. | All platforms, except Summit X450-G2, X440, and X430 | 2,000 |
| MPLS RSVP-TE transit LSPs— maximum number of transit LSPs. | All platforms, except Summit X450-G2, X440, and X430 | 2,000 |
| MPLS RSVP-TE paths— maximum number of paths. | All platforms, except Summit X450-G2, X440, X430, and X670-G2 Summit X670-G2 | 1,000 2,000 |
| MPLS RSVP-TE profiles— maximum number of profiles. | All platforms, except Summit X440, X430, X670-G2, and X450-G2 Summit X670-G2 | 1,000 2,000 |
| MPLS RSVP-TE EROs— maximum number of EROs per path. | All platforms, except Summit X450-G2, X440, and X430 | 64 |
| MPLS RSVP-TE fast reroute— MPLS RSVP-TE fast reroute (FRR) switching time. | E4G-200 | 50 msec |
| MPLS LDP peers—maximum number of MPLS LDP peers per switch. | BlackDiamond 8900 xl-series, 8900-40G6x-xm BlackDiamond X8 series E4G-400, E4G-200 Summit X460, Summit X670 Summit X670-G2, X460-G2 Summit X480, Summit X480 (40G VIM), X670V-48t, X770, X670v-48t | 64 64 32 32 128 64 |
| MPLS LDP adjacencies— maximum number of MPLS LDP adjacencies per switch. | BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series E4G-200, E4G-400 Summit X460, X480, X670, X460-G2 Summit X670V-48t, X480 (40G VIM), X770, X670-G2 | 50 64 50 50 50 64 |
| MPLS LDP ingress LSPs— maximum number of MPLS LSPs that can originate from a switch. | BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 series E4G-200 E4G-400 Summit X460, X480, Summit X670, X670V-48t, X480 (40G VIM), X770 Summit X670-G2 Summit X460-G2 | 4,000 2,048 2,048 2,048 4,000 4,000 2,048 2,048 4,000 |
| MPLS LDP-enabled interfaces— maximum number of MPLS LDP configured interfaces per switch. | Summit X460, X670 Summit X480, X670V-48t, X770 Summit X670-G2, X460-G2 BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series E4G-200, E4G-200 | 32 64 128 64 64 64 32 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|----------------|
| MPLS LDP sessions—maximum number of MPLS LDP sessions. | BlackDiamond 8900 xl-series | 64 |
| | BlackDiamond 8900-40G6x-xm | 64 |
| | BlackDiamond X8 series | 64 |
| | Summit X770, X670v-48t, X480 | 64 |
| | Summit X670-G2, X460-G2 | 128 |
| | Summit X670, X460 | 32 |
| | E4G-200, E4G-400 | 32 |
| MPLS LDP transit LSPs—maximum number of MPLS transit LSPs per switch. | BlackDiamond 8900 xl-series | 4,000 |
| | BlackDiamond 8900-40G6X-xm | 3,000 |
| | BlackDiamond X8 series | 4,000 |
| | E4G-200 | 2,700 |
| | E4G-400 | 4,000 |
| | Summit X460, X480, X770, X670V-48t, X670-G2, X460-G2 Summit X670, X480 (VIM3-40G4x) | 4,000 3,000 |
| MPLS LDP egress LSPs—maximum number of MPLS egress LSPs that can terminate on a switch. | BlackDiamond 8900 xl-series | 7,000 |
| | BlackDiamond 8900-40G6X-xm | 3,000 |
| | BlackDiamond X8 series | 7,000 |
| | E4G-200 | 2,700 |
| | E4G-400 | 6,700 |
| | Summit X460, X480, X670V-48t | 7,000 |
| | Summit X670, X480 (VIM3-40G4x) | 3,000 |
| | Summit X770 Summit X670-G2, X460-G2 | 8,000 4,000 |
| MPLS static egress LSPs—maximum number of static egress LSPs. | BlackDiamond 8900 xl-series, BlackDiamond X8 | 7,116 |
| | BlackDiamond 8900-40G | 3,020 |
| | Summit X460, X480, X670V-48t, X460-G2 | 7,116 |
| | Summit X480 (VIM3-40G4x), X670 | 3,020 |
| | Summit X770 | 8,000 |
| | Summit X670-G2 | 15,308 |
| | E4G-200 | 2,700 |
| | E4G-400 | 6,860 |
| MPLS static ingress LSPs—maximum number of static ingress LSPs. | BlackDiamond 8900 xl-series | 4,000 |
| | BlackDiamond 8900-40G | 2,048 |
| | BlackDiamond X8 | 2,048 |
| | Summit X460, X480, X460-G2 | 4,000 |
| | Summit x480-40G, X670, x670V-48t, X770, X670-G2 | 2,048 |
| | E4G-200 | 2,048 |
| | E4G-400 | 4,000 |
| MPLS static transit LSPs—maximum number of static transit LSPs | BlackDiamond 8900 xl-series | 4,000 |
| | BlackDiamond 8900-40G | 3,000 |
| | BlackDiamond X8 | 4,000 |
| | Summit X460, X480, X670V-48t, X770, X670-G2, X460-G2 | 4,000 |
| | Summit X480-40G, X670 | 3,000 |
| | E4G-200 | 2,700 |
| | E4G-400 | 4,000 |
| MSDP active peers—maximum number of active MSDP peers. | BlackDiamond 8000 series, 8900 series, X8 series | 64 |
| | Summit X460, X480, X670, E4G-400, X670-G2, X460-G2, X450-G2 | 16 |
| | Summit X770 | 64 |
| | | |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|------------------|
| MSDP SA cache entries— maximum number of entries in SA cache. | BlackDiamond 8000 series, 8900 series, X8 series | 16,000 |
| | Summit X480, X670, E4G-400 | 8,000 |
| | Summit X670-G2, X770 | 14,000 |
| | Summit X460-G2 | 10,000 |
| | Summit X450-G2 Summit X460 | 8,000 6,000 |
| MSDP maximum mesh groups —maximum number of MSDP mesh groups. | BlackDiamond 8000 series, 8900 series, X8 series | 16 |
| | Summit X460, X480, X670, E4G-400 | 4 |
| | Summit X770, X670-G2, X450-G2, X460-G2 | 16 |
| Multicast listener discovery (MLD) IPv6 multicast data sender—maximum number of IPv6 multicast streams supported on a switch. ^{im} Note: Assumes source-group- vlan mode. For additional limits, see: <ul style="list-style-type: none"> • Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode • Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode | BlackDiamond 8800 c-series | 1,000 |
| | BlackDiamond 8800 e-series | 250 |
| | BlackDiamond 8900 c-series, 8900-40G6X-xm | 1,000 |
| | BlackDiamond 8900 xl-series | 3,000 |
| | BlackDiamond X8 series | 3,000 |
| | E4G-200 | 1,500 |
| | E4G-400 | 3,000 |
| | Summit X440 | 90 |
| | Summit X460, X480 | 3,000 |
| | Summit X670 | 1,500 |
| | Summit X770, X670-G2 | 30,000 |
| | X460-G2 Summit X450-G2 | 14,000 10,000 |
| Multicast listener discovery (MLD) snooping per-VLAN filters—maximum number of VLANs supported in per-VLAN MLD snooping mode. | BlackDiamond e-series | 250 |
| | BlackDiamond 8800 c-series | 1,000 |
| | BlackDiamond 8900 c-series | 500 |
| | BlackDiamond 8900 xl-series | 2,000 |
| | BlackDiamond 8900-40G6X-xm | 500 |
| | BlackDiamond X8 a-series | 500 |
| | BlackDiamond X8 xl-series | 2,000 |
| | Summit X460, X450-G2, E4G-400 | 1,000 |
| | Summit X460-G2 | 1,200 |
| | Summit X480 | 2,000 |
| | Summit X440 | 250 |
| | Summit X670, E4G-200 | 500 |
| | Summit X770, X670-G2 | 1,200 |
| Summit X450-G2 | 512 | |
| Multicast listener discovery (MLD)v1 subscribers— maximum number of MLDv1 subscribers per port. ⁿ | BlackDiamond 8800 c-series | 500 |
| | BlackDiamond xl-series, X8 Series | 1,500 |
| | Summit X440 | 750 |
| | Summit X460, X480, X670, E4G-400 Summit X770, X670-G2, X450-G2, X460-G2 | 1,500 4,000 |
| Multicast listener discovery (MLD)v1 subscribers— maximum number of MLDv1 subscribers per switch. ⁿ | BlackDiamond 8800 series, X8 series | 10,000 |
| | Summit X440 | 5,000 |
| | Summit X460, X480, X670, E4G-400, X460-G2, X450-G2 | 10,000 |
| | Summit X770, X670-G2 | 30,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|---|
| Multicast listener discovery (MLD)v2 subscribers —maximum number of MLDv2 subscribers per port. ⁿ | BlackDiamond 8800 c-series BlackDiamond xl series BlackDiamond X8 series Summit X440, X450-G2, SummitStack Summit X460, X480, X670, E4G-400, Summit X770, X670-G2, X450-G2, X460-G2 | 500 2,500 2,000 1,000 2,000 4,000 |
| Multicast listener discovery (MLD)v2 subscribers —maximum number of MLDv2 subscribers per switch. ⁿ | BlackDiamond 8800 series, BlackDiamond xl series Summit X440, SummitStack Summit X460, X480, X670, E4G-400, X460-G2, X450-G2 Summit X770, X670-G2 | 10,000 5,000 10,000 30,000 |
| Multicast listener discovery (MLD)v2 maximum source per group —maximum number of source addresses per group. | All platforms, except Summit X430 | 200 |
| Multicast listener discovery (MLD) SSM-map entries —maximum number of MLD SSM mapping entries. | All platforms | 500 |
| Multicast listener discovery (MLD) SSM-MAP entries —maximum number of sources per group in MLD SSM mapping entries. | All platform | 50 |
| Multicast VLAN registration (MVR) —maximum number of MVR senders per switch (IP multicast compression enabled). Note: Assumes source-group-vlan mode. For additional limits, see: Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode. ^l | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X and X8 xl-series BlackDiamond 8900-40G6X-xm module Summit X440 Summit X460, E4G-400 Summit X480 Summit X670 VIM4-40G4x Summit X770, X670-G2 Summit X450-G2 Summit X460-G2 | 6,000 ⁱ 500 ⁱ 6,000 ⁱ 12,000 ^b 6,000 ⁱ 59,000 3,000 ⁱ 192 ⁱ 6,000 ⁱ 12,000 ^b 3,000 ⁱ 77,500 21,000 ^h 26,000 |
| Network login —maximum number of clients being authenticated on MAC-based VLAN enabled ports. | BlackDiamond 8000 series (clients per module/per system) BlackDiamond X8 series Summit series | 1,024 |
| Network login —maximum number of clients being authenticated with policy mode enabled. | Summit X450-G2, X460-G2 Summit X670-G2, X770 | 1,024 512 |
| Network login —maximum number of dynamic VLANs. | All platforms | 2,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|---|
| Network login VLAN VSAs —maximum number of VLANs a client can be authenticated on at any given time. | All platforms | 10 |
| ONEPolicy Roles/Profiles —maximum number of policy roles/profiles. | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | 63 N/A |
| ONEPolicy Rules per Role/Profile —maximum number of rules per role/policy. | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | Up to 952 N/A |
| ONEPolicy Authenticated Users per Switch —maximum number of authenticated users per switch. | Summit X450-G2, X460-G2 Summit X670-G2, X770 All other platforms | Up to 1,024 Up to 512 N/A |
| ONEPolicy Authenticated Users —maximum authenticated users with a combination of TCI disabled/enabled profiles. | Summit X450-G2, X460-G2 Summit X670-G2, X770 All other platforms | 682-1,022 341-510 N/A |
| ONEPolicy Authenticated Users per Port —maximum number of authenticated users per port. | Summit X450-G2, X460-G2 Summit X670-G2, X770 All other platforms | Unlimited up to 1,024 Unlimited up to 512 N/A |
| ONEPolicy Permit/Deny Traffic Classification Rules Types —total maximum number of unique permit/deny traffic classification rules types (system/stack). | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | 952 N/A |
| ONEPolicy Permit/Deny Traffic Classification Rules Types —maximum number of unique MAC permit/deny traffic classification rules types (macsource/macdest). | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | 256 N/A |
| ONEPolicy Permit/Deny Traffic Classification Rules Types —maximum number of unique IPv6 permit/deny traffic classification rules types (ipv6dest). | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | 256 N/A |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|---|
| 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). | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | 256 N/A |
| ONEPolicy Permit/Deny Traffic Classification Rules Types —maximum number of unique Layer 2 permit/deny traffic classification rules (ethertype/port). | Summit X450-G2, X460-G2, X670-G2, X770 All other platforms | 184 N/A |
| OSPFv2/v3 ECMP —maximum number of equal cost multipath OSPFv2 and OSPFv3. | All platforms, except Summit X450-G2, X440, X430, and E4G-200) BlackDiamond 8800 G48Te2 (for IPv6) Summit X450-G2 E4G-200 | 16 N/A 4 8 |
| OSPFv2 areas —as an ABR, how many OSPF areas are supported within the same switch. | All platforms (except X430, X440) | 8 |
| OSPFv2 external routes —recommended maximum number of external routes contained in an OSPF LSDB. | BlackDiamond 8000 series BlackDiamond 8900 xl-series BlackDiamond X8 series BlackDiamond X8 xl-series Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 Summit X480 E4G-200, E4G-400 | 20,000 130,000 20,000 130,000 5,000 130,000 5,000 |
| 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. | BlackDiamond 8000 series, 8900 xl-series, X8 series Summit X460, X670, X670-G2, X460-G2 Summit X480, X770 E4G-400 | 7,000 2,000 7,000 2,000 |
| OSPFv2 interfaces —recommended maximum number of OSPF interfaces on a switch (active interfaces only). | All platforms (except X430), with Advanced Edge license. All platforms (except X430 and X440) with Core license or higher | 4 400 |
| OSPFv2 links —maximum number of links in the router LSA. | All platforms, except Summit X450-G2, X770, and X430 Summit X450-G2 Summit X770 | 400 4 419 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|-----------------------------|
| OSPFv2 neighbors—maximum number of supported OSPF adjacencies. | BlackDiamond 8000 series | 128 |
| | BlackDiamond 8900 xl-series, X8 Series | 255 |
| | Summit X460, X670, X770, X440, X670-G2, X460-G2 | 128 |
| | Summit X480 | 255 |
| | Summit X450-G2 | 4 |
| OSPFv2 routers in a single area—recommended maximum number of routers in a single OSPF area. | E4G-400, E4G-200 | 128 |
| | BlackDiamond 8000 series, X8 series | 100 |
| | BlackDiamond 8900 xl-series | 200 |
| | Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 | 50 |
| | Summit X480 | 200 |
| OSPFv2 virtual links—maximum number of supported OSPF virtual links. | E4G-400 | 50 |
| | All platforms (except X450-G2, X430, and X440) with Core license or higher | 32 |
| OSPFv3 areas—as an ABR, the maximum number of supported OSPFv3 areas. | Summit X450-G2 | 4 |
| | All platforms (except X430 and X440) with Core license or higher | 16 |
| OSPFv3 external routes—recommended maximum number of external routes. | BlackDiamond 8000 series, X8 series | 10,000 |
| | BlackDiamond X8 xl-series, 8900 xl-series | 60,000 |
| | Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 | 10,000 |
| | Summit X480 | 60,000 |
| | E4G-400 | 10,000 |
| OSPFv3 inter- or intra-area routes—recommended maximum number of inter- or intra-area routes. | BlackDiamond 8000 series, 8900 xl-series, X8 series | 6,000 |
| | Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 | 3,000 |
| | Summit X480 | 6,000 |
| | E4G-400 | 3,000 |
| | OSPFv3 interfaces—maximum number of OSPFv3 interfaces. | All platforms (except X430) |
| Note: Active interfaces only, with Advanced Edge license. (See below for Core license limits.) | | |
| BlackDiamond 8000 series, BlackDiamond X8 series | | 256 |
| OSPFv3 neighbors—maximum number of OSPFv3 neighbors. | BlackDiamond 8900 xl-series | 384 |
| | Summit X460, X670, X770 | 128 |
| | Summit X480 | 384 |
| | Summit X670-G2, X460-G2 | 256 |
| | E4G-200, E4G-400 | 256 |
| | Note: With Core license or higher. (See above for Advanced Edge license limits.) | |
| | BlackDiamond 8000 series, BlackDiamond X8 series | 64 |
| BlackDiamond 8900 xl-series | 128 | |
| Summit X460, X670, X770, X670-G2, X460-G2, X450-G2 | 64 | |
| Summit X480 | 128 | |
| E4G-400 | 64 | |
| OSPFv3 virtual links—maximum number of OSPFv3 virtual links supported. | All platforms (except X450-G2, X430, and X440) with Core license or higher | 16 |
| | Summit X450-G2 | 4 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|---|
| <p>PIM IPv4 snooping—maximum number of (S,G) entries programmed in the hardware (IP multicast compression enabled).</p> <p>Note: Assumes source-group-vlan mode. For additional limits, see:</p> <ul style="list-style-type: none"> Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series BlackDiamond X8-100G4X and X8 xl-series E4G-200 E4G-400 Summit X440 Summit X480 Summit X460 Summit X670 Summit X770, X670-G2 Summit X450-G2 Summit X460-G2 | 6,000 ⁱ 500 ⁱ 6,000 ⁱ 12,000 ⁱ 3,000 ⁱ 6,000 ⁱ 59,000 ⁱ 3,000 ⁱ 6,000 ⁱ 192 ⁱ 12,000 ⁱ 6,000 ⁱ 3,000 ⁱ 77,500 21,000 26,000 |
| <p>PIM IPv4—maximum routes—maximum number of (S,G) entries installed in the hardware (IP multicast compression enabled).</p> <p>Note: Assumes source-group-vlan mode. For additional limits, see: Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode.ⁱ</p> | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series BlackDiamond X8-100G4X and X8 xl-series BlackDiamond 8900-40G6X-xm modules E4G-200 E4G-400 Summit X440 Summit X480 Summit X460 Summit X670 Summit X770, X670-G2 Summit X450-G2 Summit X460-G2 | 6,000 ⁱ 500 ⁱ 6,000 ⁱ 12,000 ⁱ 3,000 ⁱ 6,000 ^f 60,000 ⁱ 6,000 ⁱ 3,000 ⁱ 192 ⁱ 12,000 ⁱ 6,000 ⁱ 3,000 ⁱ 77,500 21,000 26,000 |
| <p>PIM IPv4-SSM (maximum SSM routes)—maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multicast compression enabled).</p> <p>Note: Assumes source-group-vlan mode. For additional limits, see: Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode.ⁱ</p> | BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond 8900-40G6X-xm BlackDiamond X8 a-series BlackDiamond X8-100G4X and X8 xl-series E4G-200 E4G-400 Summit X440 Summit X480 Summit X460 Summit X670 Summit X770, X670-G2 Summit X450-G2 Summit X460-G2 | 6,000 ⁱ 500 ⁱ 6,000 ⁱ 12,000 ⁱ 3,000 ⁱ 6,000 ⁱ 59,000 ⁱ 6,000 ⁱ 3,000 ⁱ 192 ⁱ 12,000 ⁱ 6,000 ⁱ 3,000 ⁱ 77,500 21,000 26,000 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|---------------------|
| PIM IPv6 (maximum routes) —maximum number of (S,G) entries installed in the hardware. Note: Assumes source-group-vlan mode. | BlackDiamond 8800 c-series | 1,000 |
| | BlackDiamond 8800 e-series | 250 |
| | BlackDiamond 8900 c-series | 1,000 |
| | BlackDiamond 8900-40G6X-xm | 1,000 |
| | BlackDiamond 8900 xl-series | 3,000 |
| | BlackDiamond X8 a-series | 3,000 |
| | BlackDiamond X8-100G4X and X8 xl-series | 30,000 ⁱ |
| | E4G-200 | 1,500 |
| | E4G-400 | 3,000 |
| | Summit X440 | 90 |
| | Summit X460, X480, X670 | 3,000 |
| | Summit X770, X670-G2 | 30,000 |
| | Summit X450-G2 | 10,000 |
| Summit X460-G2 | 14,000 | |
| PIM IPv4 (maximum interfaces) —maximum number of PIM active interfaces. | All platforms, except Summit X430 and X440 | 512 |
| | Summit X440 | 253 |
| PIM IPv4 (maximum interfaces) —maximum number of PIM-snooping enabled interfaces. | All platforms, except Summit X430 | 512 |
| PIM IPv4 Limits —maximum number of multicast groups per rendezvous point. | All platforms, except Summit X430 | 180 |
| PIM IPv4 Limits —maximum number of multicast sources per group. | BlackDiamond 8800 (E-series modules) | 1,000 |
| | BlackDiamond 8800 (C-series modules) | 3,000 |
| | BlackDiamond 8800 (xl-series modules) | 4,000 |
| | BlackDiamond X8 | 3,000 |
| | Summit X460-G2, X670-G2, X770, X450-G2 | 5,000 |
| | Summit X460, X480 | 1,200 |
| | Summit X670-48x | 1,000 |
| | Summit X670-48t | 4,000 |
| Summit X440 | 175 | |
| PIM IPv4 Limits —maximum number of dynamic rendezvous points per multicast group. | All platforms, except Summit X430 | 145 |
| PIM IPv4 Limits —static rendezvous points. | All platforms, except Summit X430 | 32 |
| PIM IPv6 (maximum interfaces) —maximum number of PIM active interfaces. | All platforms, except Summit X430 | 512 |
| PIM IPv6 Limits —maximum number of multicast group per rendezvous point. | All platforms, except Summit X430 | 70 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|-------|
| PIM IPv6 Limits —maximum number of multicast sources per group. | BlackDiamond 8000 | 1,280 |
| | BlackDiamond X8 | 1,500 |
| | Summit X460-G2, X670-G2 | 2,500 |
| | Summit X460, X480 | 800 |
| | Summit X670 | 2,000 |
| | Summit X440 | 175 |
| | Summit X450-G2 | 2,000 |
| | Summit X770 | 2,500 |
| PIM IPv6 Limits —maximum number of dynamic rendezvous points per multicast group. | All platforms, except Summit X430 | 64 |
| PIM IPv6 Limits —maximum number of secondary address per interface. | All platforms, except Summit X430 | 70 |
| PIM IPv6 Limits —static rendezvous points. | All platforms, except the Summit X430 | 32 |
| Policy-based routing (PBR) redundancy —maximum number of flow-redirects. | All platforms | 256° |
| Policy-based routing (PBR) redundancy —maximum number of next hops per each flow-direct. | All platforms | 32° |
| Port-specific VLAN tags —maximum number of port-specific VLAN tags. | All platforms, except Summit X450-G2, X440, and X430 | 1,023 |
| Port-specific VLAN tags —maximum number of port-specific VLAN tag ports. | BlackDiamond X8 and 8800 xl-series | 8,090 |
| | Summit X480 | 3,800 |
| | Summit X460-48t | 7,200 |
| | Summit X460-24x, X670-48x | 3,400 |
| | Summit X670V-48t | 3,600 |
| | Summit X670v-48t stack | 7,200 |
| | Summit X770, X670-G2 | 6,400 |
| | Summit X460-G2 | 4,000 |
| | E4G-400 | 3,400 |
| E4G-200 | 3,800 | |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|--|
| Private VLANs —maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN. | BlackDiamond 8800 c-, e-, xl-series with eight modules of 48 ports 8900-G96T-c modules | 383 |
| | BlackDiamond X8 series Summit X770 Summit X670-G2, X670v-48t Summit X670 Summit X480 Summit X460-G2, X460 Summit X440 Summit X430 Summit X450-G2 E4G-200 E4G-400 | 767 103 63 47 23 53 25 27 51 11 33 |
| 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. | Summit X770, X670-G2, X460-G2, X450-G2 Summit X670, X480, X460, X460, X480 Summit X440 E4G-200, E4G-400 | 1,024 512 127 512 |
| Private VLANs —maximum number of private VLANs in an L2-only environment. | BlackDiamond 8800 c-, e-series BlackDiamond 8900 series, X8 series E4G-200 E4G-400 Summit X440 Summit X480, Summit X670 Summit X460 Summit X770, X670-G2, X460-G2, X450-G2 Summit X430 | 384 2,046 597 1,280 127 597 820 1,280 255 |
| PTP/1588v2 Clock Ports | Summit X770, X460-G2, X670-G2, and E4G-200, E4G-400 cell site routers | 32 for boundary clock 1 for ordinary clock |
| PTP/1588v2 Clock Instances | Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers | 2 combinations: <ul style="list-style-type: none"> • Transparent clock + ordinary clock • Transparent clock + boundary clock |
| PTP/1588v2 Unicast Static Slaves | Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers | 40 entries per clock port |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|---|
| PTP/1588v2 Unicast Static Masters | Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers | 10 entries per clock type |
| 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. | All platforms, except Summit X430 | 10,000 |
| RIP neighbors —maximum number of RIP neighbors. | E4G-200 | 256 |
| RIP interfaces on a single router —recommended maximum number of RIP routed interfaces on a switch. | BlackDiamond 8000 series, X8 series BlackDiamond 8900 xl-series Summit X440 Summit X460, X670-G2, X460-G2 Summit X480 Summit X670, X770, X450-G2 E4G-400 | 256 384 128 256 384 256 256 |
| RIPng learned routes —maximum number of RIPng routes. | BlackDiamond 8000 series, X8 series BlackDiamond 8900 xl-series Summit X480 Summit X460, X670, X670-G2, X460-G2, X770, X450-G2 E4G-200 | 3,000 5,000 5,000 3,000 3,000 |
| Spanning Tree (maximum STPDs) —maximum number of Spanning Tree Domains on port mode EMISTP. | All platforms (except Summit X430, X440) Summit X440 Summit X430 | 64 32 16 |
| Spanning Tree PVST+ —maximum number of port mode PVST domains. Note: <ul style="list-style-type: none"> Maximum of 10 active ports per PVST domain when 256 PVST domains are configured. Maximum of 7 active ports per PVST domain when 128 PVST domains are configured. | BlackDiamond X8 and 8900 series Summit X670, X770, X670-G2 Summit X460, X480, X440, X460-G2 Summit X430 Summit X450-G2 E4G-400 | 256 256 128 50 128 128 |
| Spanning Tree —maximum number of multiple spanning tree instances (MSTI) domains. | All platforms (except Summit X430, X440) Summit X430 Summit X440 | 64 5 32 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|-------|
| 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. | BlackDiamond X8, 8800, 8900 MSM 128/XL | 500 |
| | Summit X770, X670-G2, X670v-48t, X670 | 500 |
| | Summit X480, X460-G2, X460, X450-G2 | 600 |
| | E4G-200 | 500 |
| | E4G-400 | 600 |
| | Summit X440 | 250 |
| | Summit X430 | 100 |
| Spanning Tree —maximum number of VLANs on all MSTP instances. | BlackDiamond X8, 8800, 8900 MSM 128/XL | 1,000 |
| | Summit X770 | 1,024 |
| | Summit X670-G2, X670v-48t, X670, X480 | 1,000 |
| | Summit X460-G2, X460, X450-G2 | 1,024 |
| | E4G-200 | 1,000 |
| | E4G-400 | 1,024 |
| | Summit X440 | 500 |
| | Summit X430 | 200 |
| 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. | All platforms (except Summit X430, X440) | 4,096 |
| | Summit X430 | 1,024 |
| | Summit X440 | 2,048 |
| Spanning Tree (maximum VLANs) —maximum number of STP-protected VLANs (dot1d and dot1w). | BlackDiamond X8, 8800, 8900 MSM 128/XL | 1,024 |
| | Summit X770 | 1,024 |
| | Summit X670-G2, X670v-48t, X670, X480 | 560 |
| | Summit X460-G2, X460, X450-G2 | 600 |
| | E4G-200 | 500 |
| | E4G-400 | 600 |
| | Summit X440 | 500 |
| | Summit X430 | 128 |
| SSH (number of sessions) —maximum number of simultaneous SSH sessions. | All platforms | 8 |
| Static MAC multicast FDB entries —maximum number of permanent multicast MAC entries configured into the FDB. | BlackDiamond 8000 c-, e-, xl-series | 1,024 |
| | BlackDiamond X8 series | |
| | All Summits | |
| | E4G-200, E4G-400 | |
| Syslog servers —maximum number of simultaneous syslog servers that are supported. | All platforms | 4 |
| Telnet (number of sessions) —maximum number of simultaneous Telnet sessions. | All platforms | 8 |
| TRILL —trees rooted from switch. | BlackDiamond X8 Summit X670, X770 | 1 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|--|-------------------------------|
| TRILL—computed trees. | BlackDiamond X8 Summit X670, X770 | 1 |
| TRILL—TRILL VLANs. | BlackDiamond X8 Summit X670, X770 | 4 |
| TRILL—forwarding VLANs. | BlackDiamond X8 Summit X670, X770 | 4,095 |
| TRILL—forwarding ports. | BlackDiamond X8 Summit X670, X770 | All |
| TRILL—RBridge FDB entries. | BlackDiamond X8 Summit X670 Summit X770 | 128,000 128,000 288,000 |
| TRILL—ECMP RBridge next hops. | BlackDiamond X8 Summit X670, X770 | 8 |
| TRILL—neighbor adjacencies. | BlackDiamond X8 Summit X670, X770 | 32 |
| TRILL—nodes. | BlackDiamond X8 Summit X670, X770 | 256 |
| TRILL—links. | BlackDiamond X8 Summit X670, X770 | 2,000 |
| Virtual routers —maximum number of user-created virtual routers that can be created on a switch. Note: Virtual routers are not supported on Summit X440 series switches. | BlackDiamond 8000 c-, xl-, xm-series BlackDiamond X8 series Summit X460, X460-G2, X480, X670, X670-G2, X770, X450-G2 E4G-200, E4G-400 | 63 |
| Virtual router forwarding (VRFs) —maximum number of VRFs that can be created on a switch. Note: * Subject to other system limitations. | All platforms, except Summit X440 and X430 | 960 * |
| Virtual router protocols per VR —maximum number of routing protocols per VR. | All platforms, except Summit X440, X430 | 8 |
| Virtual router protocols per switch —maximum number of VR protocols per switch. | All platforms, except Summit X440, X430 | 64 |
| VLAN aggregation —maximum number of port-VLAN combinations on any one superVLAN and all of its subVLANs. | All platforms (except Summit X430, X440) Summit X440, X430 | 1,000 256 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|---|
| VLANS —includes all VLANs. Note: ExtremeXOS supports only 4,092 user-configurable VLANs. (VLAN 1 is the default VLAN, and 4,095 is the management VLAN, and you may not configure them.) | All platforms | 4,094 |
| VLANS —maximum number of port-specific tag VLANs. | BlackDiamond 8800 xl-series only, BlackDiamond X8 series BlackDiamond X8 xl-series Summit X460, X770, X480, E4G-400, X670-G2, X460-G2 Summit X670, X670V-48t E4G-400 E4G-200 | 1,023 4,093 4,093 1,023 4,093 2,047 |
| VLANS —maximum number of port-specific tag VLAN ports. | BlackDiamond 8800 xl-series only BlackDiamond X8 BlackDiamond X8 xl-series E4G-400, E4G-200 Summit X460, X670, X670V-48t, X460-G2 Summit X770, X670-G2 Summit X480 | 4096 4096 32,767 4096 4096 8,192 16,383 |
| VLANS (Layer 2) —maximum number of Layer 2 VLANs. | All platforms | 4,094 |
| VLANS (Layer 3) —maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs. | BlackDiamond X8 Summit X460-G2, X670, X770, X670-G2, X450-G2 Summit X440 Summit X480, X460 E4G-200, E4G-400 | 2,048 2,048 254 512 512 |
| VLANS (maximum active port-based) —maximum active ports per VLAN when 4,094 VLANs are configured with default license. | BlackDiamond X8, 8800 series Summit X770, X670-G2, X670v-48t, X670, X480, X460-G2, X460, X450-G2 E4G-200 E4G-400 Summit X440 Summit X430 | 32 32 12 32 7 2 |
| VLANS (maximum active protocol-sensitive filters) —number of simultaneously active protocol filters in the switch. | All platforms, except Summit X450-G2 Summit X450-G2 | 15 16 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|---|
| VLAN translation—maximum number of translation VLANs. Assumes a minimum of one port per translation and member VLAN. | BlackDiamond 8000 a-, c-, e-, xl series with eight modules of 48 ports 8900-G96T-c modules | 383 767 |
| | Summit X770 | 103 |
| | Summit X670-G2, X670v-48t | 63 |
| | Summit X670 | 47 |
| | Summit X480 | 53 |
| | Summit X460-G2 | 53 |
| | Summit X460 | 57 |
| | E4G-200 | 11 |
| | E4G-400 | 33 |
| | Summit X440 | 25 |
| Summit X430 | 27 | |
| Summit X450-G2 | 51 | |
| 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 has tagged and translated ports. | Summit X770, X670-G2, X450-G2 Summit X670, X480, X460 Summit X440 E4G-200, E4G-400 | 1,024 512 127 512 |
| VLAN translation—maximum number of translation VLAN pairs in an L2-only environment. | BlackDiamond 8800 c-, e-series BlackDiamond 8900 xl-series BlackDiamond X8 series Summit X460 Summit X430 Summit X480, X670, X770, X670-G2, X460-G2 Summit X450-G2 Summit X440 E4G-400, E4G-200 | 384 2,046 2,046 2,000 512 2,046 1,024 127 2,000 |
| VRRP (v2/v3-IPv4) (maximum instances)—maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. | BlackDiamond X8 BlackDiamond 8800 MSM-48c and MSM-96 BlackDiamond 8900 xl-series 8900-MSM128 Summit X770, X670, X670-G2, X460-G2, X480, X450-G2 Summit X460 Summit X440 E4G-200, E4G-400 | 511 511 511 511 255 32 128 |
| VRRP (v3-IPv6) (maximum instances)—maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. (VRRP-VRRPv3-IPv6) | BlackDiamond X8, BlackDiamond 8800 MSM-48c and MSM-96 BlackDiamond 8900 xl-series 8900-MSM128 Summit X770, X670, X670-G2, X460-G2, X450-G2 Summit X460, X480 Summit X440 E4G-200, E4G-400 | 511 511 511 511 255 15 255 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|--|--|---|
| VRRP (v2/v3-IPv4/IPv6) (maximum VRID)—maximum number of unique VRID numbers per switch. | All platforms with Advanced Edge license or higher, except Summit X430 | 7 |
| VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN)—maximum number of VRIDs per VLAN. | All platforms with Advanced license or higher, except for Summit X430 | 7 |
| VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks)—maximum number of ping tracks per VLAN. | All platforms with Advanced Edge license or higher, except Summit X430 | 8 |
| VRRP (maximum ping tracks)—maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher. | All platforms, except the Summit X440 | 8 (20 centisecond or 1 second hello interval) |
| | Summit X440 Hello interval: 20 centiseconds Hello interval: 1 second | 1 |
| VRRP (v3-IPv6) (maximum ping tracks)—maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher. | All platforms, except the Summit X440 | 8 (20 centisecond or 1 second hello interval) |
| | Summit X440 Hello interval: 20 centiseconds Hello interval: 1 second | 1 |
| VRRP (v2/v3-IPv4/IPv6) (maximum iproute tracks)—maximum number of IP route tracks per VLAN. | All platforms with Advanced Edge license or higher, except Summit X430 | 8 |
| VRRP (v2/v3-IPv4/IPv6)—maximum number of VLAN tracks per VLAN. | All platforms with Advanced Edge license or higher, except Summit X430 | 8 |

Table 4: Supported Limits (continued)

| Metric | Product | Limit |
|---|---|-----------------------------|
| 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. | BlackDiamond 8800 c-series with 100 DACLs with 500 DACLs | 10 3 |
| | BlackDiamond 8800 c-series with 100 DACLs with 500 DACLs | 10 3 |
| | Summit X480, X670 with 100 DACLs with 500 DACLs | 4 1 |
| | Summit X450-G2 with 100 DACLs | 10 |
| XNV authentication —maximum number of VMs that can be processed (combination of local and network VMs). | All platforms, except Summit X430 and X450-G2 Summit X450-G2 | 2,048 1,024 |
| XNV database entries —maximum number of VM database entries (combination of local and network VMs). | All platforms, except Summit X430 | 16,000 |
| XNV database entries —maximum number of VPP database entries (combination of local and network VPPs). | All platforms, except Summit X430 | 2,048 |
| XNV dynamic VLAN —Maximum number of dynamic VLANs created (from VPPs / local VMs). | All platforms, except Summit X430 | 2,048 |
| XNV local VPPs —maximum number of XNV local VPPs. | All platforms, except Summit X430 | 2,048 ingress 512 egress |
| XNV policies/dynamic ACLs —maximum number of policies/dynamic ACLs that can be configured per VPP. | All platforms, except Summit X430 | 8 ingress 4 egress |
| XNV network VPPs —maximum number of XNV network VPPs. ^p | All platforms, except Summit X430 | 2,048 ingress 512 egress |

^a The table shows the total available.

^b Limit depends on setting configured for "configure forwarding external-tables".

^c When there are BFD sessions with minimal timer, sessions with default timer should not be used.

^d Based on in "none more-l2" mode.

^e Based on forwarding internal table configuration "more l2".

-
- ^f Effective capacity varies based on actual MAC addresses and VLAN IDs used and hash algorithm selected.
 - ^g Based on "I2-only mode".
 - ^h Based on forwarding internal table configuration "more I3-and-ipmc".
 - ⁱ Based on forwarding external table configuration "I3-only ipv4".
 - ^j The limit depends on setting configured with configure iproute reserved-entries.
 - ^k Based on forwarding external table configuration "I3-only ipv4".
 - ^l Based on forwarding external table configuration "I3-only ipv6".
 - ^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 would be appropriately lessened.
 - ^o Sum 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.

3 Open Issues, Known Behaviors, and Resolved Issues

Open Issues

Known Behaviors

Resolved Issues in ExtremeXOS 16.1.3-Patch1-7

Resolved Issues in ExtremeXOS 16.1.3-Patch1-4

Resolved Issues in ExtremeXOS 16.1.3-Patch1-2

Resolved Issues in ExtremeXOS 16.1.3

Resolved Issues in ExtremeXOS 16.1.2-Patch1-4

Resolved Issues in ExtremeXOS 16.1.2-Patch1-1

Resolved Issues in ExtremeXOS 16.1.2

Resolved Issues in ExtremeXOS 16.1

This chapter describes items needing further clarification and behaviors that might not be intuitive. It also includes the items that have been resolved.

Open Issues

The following are new open issues for supported features found in ExtremeXOS 16.1.3-Patch1-7.

Table 5: Open Issues, Platform-Specific, and Feature Change Requests (CRs)

| CR Number | Description |
|----------------|--|
| General | |
| xos0061053 | ExtremeXOS supports the use of RC4 in one or more cipher suites. The RC4 cipher is flawed in its generation of a pseudo-random stream of bytes so that a wide variety of small biases are introduced into the stream, decreasing its randomness. If plaintext is repeatedly encrypted (e.g., HTTP cookies), and an attacker is able to obtain many (i.e., tens of millions) ciphertexts, the attacker may be able to derive the plaintext. |
| xos0061052 | ExtremeXOS accepts connections encrypted using SSL 2.0 and/or SSL 3.0, which reportedly suffer from several cryptographic flaws. An attacker may be able to exploit these issues to conduct man-in-the-middle attacks or decrypt communications between the affected service and clients. |
| xos0060993 | Nessus scan detects the following medium vulnerabilities in ExtremeXOS: <ul style="list-style-type: none">SSH: CBC Mode Ciphers EnabledSSH: Weak Mac Algorithms Enabled |
| xos0060930 | When ONEPolicy is enabled and you reach the configured maximum number of authenticated sessions, sessions continue to attempt to authenticate, and then terminate if successful. |

**Table 5: Open Issues, Platform-Specific, and Feature Change Requests (CRs)
(continued)**

| CR Number | Description |
|--|---|
| xos0061492 | For the Summit X430 series switches, you can only create around 3,900 VLANs, which is short of the limit of 4,094. For Summit X440 series switches, you can only create 4,094 VLANs and 40–43K VPIF, whereas 53K VPIF was obtainable in ExtremeXOS 15.7.1. |
| xos0063245 | With IGMP per-VLAN mode, VRRP flaps occur after adding tagged ports to VLANs. |
| BlackDiamond 8800 Series Switches | |
| xos0060136 | With NetLogin with MAC enabled and with dynamic VLAN configured, if FDB ageout timer is configured as 50, sometimes FDB does not synchronize and the command <code>show netlogin mac</code> shows clients authenticated on nvlan itself. |
| Summit X450-G2 Series Switches | |
| xos0061097 | On Summit X450G2 stack of eight, back-to-back failovers while sending slow-path traffic across eight slots, produce the following error: <code>04/01/2015 13:36:33.65 <Error:Kern.Card.Error> Slot-5: bcm_tx_list() returned -4: Invalid parameter</code> Issue does not occur, if slow-path traffic is stopped. |
| Summit X670 Series Switches | |
| xos0062312 | On Summit X670V-48x-VIM4-40G4X switches, when you disable ports on a peer switch, additional 40G ports may go down. Note: Configuring the debounce timer to 4 seconds on these ports may resolve this issue. |
| Summit Series Switches | |
| xos0060283 | The SMON MIB (RFC 2613) which was used to configure mirroring using SNMP is not available in ExtremeXOS. |
| ACLs | |
| xos0061183 | On BlackDiamond X8 and 8800 series switches, if failover occurs during an active ESVT test, sometimes it might persist in "running" state. |
| BGP | |
| xos0060352 | BGP speaker accepts invalid updates (for example, invalid IP addresses such as 0.0.0.0/24). These are installed in BGP LOCAL RIB, as well as in route table. |
| Clocking (1588v2) | |
| xos0060785 | Precision time feature limitations for ExtremeXOS 16.1: <ul style="list-style-type: none"> • ExtremeXOS 16.1 slave ports sync to grandmasters, such as Symmetricom, and to other ExtremeXOS 16.1 clocks, but not to ExtrememXOS 15.7, and earlier. If networks of clocks are to be upgraded to ExtremeXOS 16.1, complete the upgrades simultaneously or staged starting closest to the grandmaster. Before beginning a staged upgrade, where an earlier version of ExtremeXOS must sync to an ExtremeXOS 16.1 clock, test the particular configuration beforehand. • ExtremeXOS 16.1 slave clock ports must be configured with the "slave-only" option to sync to other ExtremeXOS 16.1 clocks. |
| MPLS | |

Table 5: Open Issues, Platform-Specific, and Feature Change Requests (CRs) (continued)

| CR Number | Description |
|-----------------|---|
| xos0061018 | After failover, traffic fails across VPLS configured with 64 LSPs across LAG. |
| xos0061276 | MPLS LSP (LDP/RSVP) is not formed when BGP is used as IGP routing protocol. |
| xos0061374 | With an L2VPN session between two Label Edge Routers (LERs), broadcast packets egressing the LERs are corrupted. |
| xos0062314 | Detour LSP counters display incorrect values in the output of the commands <code>show mpls rsvp-te lsp</code> and <code>show mpls rsvp-te lsp fast-reroute</code> . |
| NetLogin | |
| xos0060488 | With upload and download of NetLogin with UPM XSF file, UPM profile is not executed for the user-authenticate and unauthenticate events. |
| xos0060280 | Enabling NetLogin mac on mirrored ports does not produce an error. |
| xos0061375 | Re-authentication fails for some NetLogin authenticated clients after changing the EXTREME_NETLOGIN_EXTENDED_VLAN VSA (211) with scaled number of NetLogin authenticated clients. |
| xos0061116 | After disabling NetLogin dot1x, attempting to enable NetLogin dot1x produces an error indicating that NetLogin is already enabled on a port. |
| OSPF | |
| xos0061100 | CPU utilization monitor incorrectly displays 99% CPU usage for OSPF while restarting OSPF process. |

Known Behaviors

The following are limitations in ExtremeXOS system architecture that have yet to be resolved.

Table 6: Known Issues, Platform-Specific, and Feature Change Requests (CRs)

| CR Number | Description |
|---------------------------------------|--|
| Summit X460-G2 Series Switches | |
| xos0059693 | Only 'macdest', 'macsource', or 'port' policy rules can be applied to QinQ (that is, double-tagged) packets received on an untagged VMAN port. |
| ACL | |
| xos0060980 | Two-stage ACL with tunnel configuration does not work when class-id is greater than "1". |
| NetLogin | |
| xos0061484 | Ports added using command line are removed by NetLogin on multiple VLAN VSA movement. |
| xos0060216 | NetLogin MAC client does not move to service unavailable VLAN with dot1x MAC-and web-enabled on same port with auth database order as local. |
| xos0060140 | Movement from non-NetLogin VLAN to multiple untagged VLANs does not occur in MAC-based mode. |

Table 6: Known Issues, Platform-Specific, and Feature Change Requests (CRs) (continued)

| CR Number | Description |
|----------------------|---|
| xos0060351 | In NetLogin web non-policy mode, you are unable to add more than one VLAN with VSA 211 (untagged VLAN) in MAC-based mode. |
| STP/RSTP/MSTP | |
| xos0058362 | Ports configured as auto with auto-edge feature turned on, do not have this status correctly shown in <code>show stpd port</code> command. Port operation mode appears as "Point-point". Note: Port operation mode appears correctly in <code>show stpd port detail</code> command. |

Resolved Issues in ExtremeXOS 16.1.3-Patch1-7

The following issues were resolved in ExtremeXOS 16.1.3-Patch1-7. ExtremeXOS 16.1.3-Patch1-7 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, ExtremeXOS 16.1.2, and ExtremeXOS 16.1.3. For information about those fixes, see the release notes for the specific release.

Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-7

| CR Number | Description |
|----------------|---|
| General | |
| xos0056729 | After being deleted from the switch, old configuration files are incorrectly visible in the help menu. |
| xos0057931 | After rebooting the switch multiple times, following error log message appears: <code><Error:cm.loadErr> Failed to load configuration: timed out (after 150 seconds) while waiting for all applications to get ready to load configuration on OPERATIONAL (eaps is still not ready yet).</code> |
| xos0061597 | After authenticating a NetLogin client, executing the command <code>clear netlogin state port</code> , and then checking that ID-mgmt has deleted the clients, log displays UnDscvrId EMS message, which does not indicate the identity user. |
| xos0062265 | Some legacy commands are not recognized. |
| xos0062720 | Unable to save configuration when ACL/CFM is configured on multiple VLANs. |
| xos0063332 | Configuration changes to VPLS are not fully retrieved by SNMP walk, which returns values for only few VPLS index. |

Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-7 (continued)

| CR Number | Description |
|---------------------------------------|--|
| xos0063980 | EDP process crash with signal 11 when receiving CDP packet with IPv6 address type. |
| xos0064009 | MLAG+EAPS:Traffic forwarding stops after EAPS that include ISC link converges. |
| xos0064054 | SNMPwalk on extremeAclStatsTable returns value with port instance instead of ifIndex. |
| xos0064055 | Resiliency Enhancement for IPv4 and IPv6 Static Routes feature has been added. |
| xos0064151 | Errors occur when removing DHCP configuration from VLANs when LAG ports are added to the VLANs. |
| xos0064223 | Need to add an ACL match condition for matching next-hop addresses during the look-up cycle of a packet, so that actions can be taken based on the next-hop a packet is destined for. |
| xos0064278 | In a SummitStack or BlackDiamond chassis, FDB is not programmed in hardware after three failovers and fallback. |
| xos0064043 | Unable to use a configuration file that has been copied from an existing configuration file. |
| xos0064129 | Policy refresh never completes with network-zone configuration. |
| xos0064215 | The following log message appears when a subnet is reachable both using MPLS and non-MPLS: <Warn:Kern.IPv4FIB.Warning> Slot-4: dest 0x0A420000 / 24 nexthop 0xAC11121E: Unable to add route to unit 1, rc Entry not found. Shadow problem |
| xos0064307 | RADIUS accounting configuration is incorrect as shown by the command <code>show conf aaa</code> and is lost after upgrade. |
| Summit X430 Series Switches | |
| xos0064084 | In Summit X430 series switches, the command <code>show power details</code> displays fan status as "empty". |
| Summit X440 Series Switches | |
| xos0063627 | ARP is not re-added to hardware after it is removed initially due to the table being full. |
| xos0064049 | Need command provision to tune pre-emphasis settings for stacking port on Summit X440 series switches. |
| xos0064050 | While running diagnostics on a Summit X440 10G model switch with revision 10 and diagnostics test version 6.0 or above, "Test loopback phy fiber" and "Test snake interface" fail. |
| Summit X460-G2 Series Switches | |

Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-7 (continued)

| CR Number | Description |
|---------------------------------|---|
| xos0063811 | Summit X460-G2 series switches with ExtremeXOS 15.6 through 21.1, have the following limitations for SyncE input reference frequency: <ul style="list-style-type: none"> Network clock does not lock with input SyncE source port 52 (both at 10G and 1G speed) on all 48-port models (X460-G2-48t, 48x and 48p). <p>Workaround: For SyncE input at 10G, avoid port 52.</p> <ul style="list-style-type: none"> When the 10G ports operate at 1G speed, the network clock does not lock. Models with Ethernet BASE-T or Power over Ethernet (PoE) ports may lock on initial configuration, but do not lock after a save and reboot. <p>Workaround: For SyncE input at 1G, use a 1G port, not a 10G port.</p> |
| SummitStack | |
| xos0061834 | In SummitStacks, the command <code>synchronize stacking slot slot no</code> does not work from master node if the target slot is in failed state. |
| xos0061861 | A per-port meter configured on a SummitStack may not be properly configured on the backup node following a reboot. |
| BlackDiamond X8 Switches | |
| xos0062262 | UDP profile configurations on VLANs do not take effect in BlackDiamond X8 series switches. |
| xos0063080 | ACL rule with destination zone match condition does not take effect, since IP address gets installed in reverse order in hardware. |
| xos0063958 | BDX8-40G12X-XL module goes into reboot cycle if any physical loop is configured or a network loop is present on the module. |
| xos0064010 | The command <code>show port buffer</code> displays an incorrect port range for 100G I/O modules. |
| xos0064201 | On BlackDiamond X8 chassis with BDXA-40G24X modules loaded in all I/O slots, some modules may occasionally time out while booting up. |
| xos0064444 | Management module crashes continuously when there is a physical loop on the front panel ports of the I/O module. |

Resolved Issues in ExtremeXOS 16.1.3-Patch1-4

The following issues were resolved in ExtremeXOS 16.1.3-Patch1-4. ExtremeXOS 16.1.3-Patch1-4 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, ExtremeXOS 16.1.2, and ExtremeXOS 16.1.3. For information about those fixes, see the release notes for the specific release.

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-4

| CR Number | Description |
|--|--|
| General | |
| xos0056368 | Kernel errors occur after disabling sharing configuration on ISC ports of MLAG. For example: "exvlan: handleVsmKernelRequest:8545: handleVsmKernelRequest Invalid Ingress port: 1000008 got" |
| xos0061841 | FDB entries are not learned again after limit learning is unconfigured, and then configured again, with PSTAG configuration in SummitStacks. |
| xos0063082 | Updating DSCP value is not getting refreshed for Dynamic ACL. |
| xos0063463 | Static FDB created on PStag VLAN port is incorrectly displayed in show configuration. |
| xos0063761 | Traffic not getting forwarded after disable/enable sharing when traffic ingressing port is part of both port specific tag (PStag) and service VMAN (untagged port). |
| xos0063948 | Clearflow delta values are randomly not calculated properly. |
| xos0063968 | HAL process ends unexpectedly after changing/reverting service VLAN tag. |
| xos0064063 | Packet-Resolution match conditions need to be added as ACL match conditions. |
| xos0064067 | Traffic loss occurs in MLAG setup when ingress port and ISC port reside on different hardware unit and when internal port number for both these ports are same. |
| xos0064179 | MAC movement occurs in switch acting as an STP root bridge when PVST + BPDUs are sent by peer switch using STP blocked port. |
| BlackDiamond 8800 Series Switches | |
| xos0063511 | Memory allocation failure messages appear with MSM-96 after failover: "MSM-B: RX: Failed to allocate mem, unit 0, chan 1 dv 0, i 15, rv -2" |
| BlackDiamond X8 Switches | |
| xos0063546 | The following error message may appear when BlackDiamond X8 modules are rebooted after configuring port partition: Erro:Kern.Card.Error> Slot-2: _setSchedMode: u:p=1:006 schedMode=1 err=Invalid parameter |
| Summit X670 Series Switches | |
| xos0063263 | On Summit X670 series switches, 1000BaseSX optics are incorrectly detected as 100BaseFX. |
| Summit Family Switches | |
| xos0059831 | In Summit X670 and X770 series switches, the QSFP+ end of 10GB-4-F10-QSFP cable is incorrectly detected as non-Extreme certified. |
| xos0061506 | In Summit X440-G2 and X460-G2 series switches, the combo port comes up as active even though when link peer port is down. |
| xos0062700 | When upgrading from ExtremeXOS 15.7 or earlier to 16.1, image download fails if image was installed in backup node first and master node second. |

Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-4 (continued)

| CR Number | Description |
|------------|--|
| xos0063548 | LX100 optics are not working on 10G ports. |
| xos0063919 | On standby nodes, IP ARP refresh and Neighbor refresh are now disabled on VR-Mgmt. Primary and backup nodes use the configured enabled/disabled setting. |

Resolved Issues in ExtremeXOS 16.1.3-Patch1-2

The following issues were resolved in ExtremeXOS 16.1.3-Patch1-2. ExtremeXOS 16.1.3-Patch1-2 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, ExtremeXOS 16.1.2, and ExtremeXOS 16.1.3. For information about those fixes, see the release notes for the specific release.

Table 9: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-2

| CR Number | Description |
|----------------|--|
| General | |
| xos0057785 | STP domain tag is removed when all ports are deleted from STP auto-bind enabled-VLANs. |
| xos0061745 | Ampersand used in UPM script is replaced by "& amp" in the XSF configuration. |
| xos0062850 | When upgrading ExtremeXOS to 15.7 or later releases, the web HTTP access is enabled even though it is disabled in the configuration. |
| xos0063134 | Traffic stops after disabling, and then enabling LAG porst having pstag with static FDB. |
| xos0063173 | Process dcbgp ends unexpectedly with signal 11 after issuing the command <code>show bgp neighbor</code> . |
| xos0063186 | Kernel oops occurs when deleting private VLAN. |
| xos0063240 | ACL process ends unexpectedly when switch has clear-flow ACL rule with count interval greater than snmptrap generation timer. |
| xos0063245 | With IGMP per-VLAN mode, VRRP flaps occur after adding tagged ports to VLANs. |
| xos0063258 | OSPFv3 routes are not updated in routing tables if there is a delay in receiving LSAs. |
| xos0063282 | ExtremeXOS CLI restricts PVLAN subscriber VLAN from being configured as an EAPS-protected VLAN. |
| xos0063418 | "No mapping for Modid" errors occur when sFlow is enabled on the port. |

Table 9: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-2 (continued)

| CR Number | Description |
|------------|---|
| xos0063423 | Memory leak occurs in ISIS process when exporting OSPF routes to ISIS. |
| xos0063429 | The output of the command <code>show fan</code> shows the fan status as empty after a hot re-seating of the fan module. |
| xos0063457 | Configuration for adding network VLAN port in STP for subscriber is not saved. |
| xos0063465 | Cannot add/delete ports to load-shared MLAG ports without disabling MLAG. |
| xos0063478 | Traffic drop occurs while adding new member port to the existing LAG group and PSTAG is configured on the port. |
| xos0063484 | Enhancement added in STP flush generation mechanism to reduce hardware programming load. |
| xos0063493 | Traffic stops after disabling, and then enabling LAG ports in VPLS service VLAN with static FDB. |
| xos0063585 | AAA process ends unexpectedly when NetLogin primary RADIUS server is down and "Active directory domain services" alone is down in secondary NetLogin RADIUS server. |
| xos0063673 | If one MAC address learned on the same port in a dynamic VLAN ages out, the whole dynamic VLAN is deleted after 20 seconds even though other MAC addresses are still in the FDB table. |
| xos0063710 | Kernel oops occurs on switch with Private VLAN and MLAG configuration after executing <code>restart ports all</code> . |
| xos0063736 | In Syslog, username information appears as "*****" during login/logout cases. |
| xos0063815 | Traffic is dropped on one VPLS instance when another VPLS instance running on same switch is deleted. |
| xos0063853 | Potential memory corruption when MAC locking is configured. |
| xos0063870 | Kernel oops occurs due to memory overrun in user kernel interface. |
| xos0063956 | ACL slice is not freed up after changing IGMP snooping filter from per-vlan to per-port mode. |
| xos0054714 | When ACLs are applied in both ingress and egress directions, you cannot see egress direction using SNMP. When a policy has more than one counter, using SNMP, you can only check the updates from the first counter, and subsequent counters do not appear. |
| xos0059481 | Static FDB is programmed incorrectly in hardware after a stack failover. |
| xos0062902 | Kernel oops occurs while disabling service VMAN with CEP configuration. |
| xos0063521 | A few IBGP routes are not updated in routing table when <code>disable bgp</code> and <code>enable bgp</code> commands are executed in quick succession. |
| xos0063547 | Process ACL ends unexpectedly after applying a policy file with source zone as a match condition. |
| xos0063644 | exsshd process ends unexpectedly with signal 11 in rare scenarios. |

Table 9: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-2 (continued)

| CR Number | Description |
|--|--|
| xos0063814 | UPM process ends unexpectedly with Signal 11 occasionally when UPM timers are configured. |
| xos0063365 | Frequent MLAG bulk syncs observed due to checksum mismatch between MLAG peers when ISC port was added as an untagged port to a tagged VLAN and VRRP was running between the peers. |
| BlackDiamond 8800 Series Switches | |
| xos0063552 | Backup Master Switch Fabric Module (MSM) fails to reach 'In sync' state after bootup in fully loaded chassis with a large number of ACL rules applied. |
| xos0063614 | Kernel crash occurs when receiving DHCP packets with invalid field values. |
| xos0063872 | After multiple executions of <code>run failover</code> with redirect-flow configuration, IPv4 ping fails. |
| xos0063333 | In BlackDiamond 8800 series switches, optics information is not detected and ports remain in "Ready" state after reboot. |
| BlackDiamond X8 Series Switches | |
| xos0057363 | Sample message with invalid slot ID received after enabling sFlow. |
| xos0063928 | In BlackDiamond X8 series switches, Sysuptime in sFlow packets is invalid. |
| E4G 400 Cell Site Routers | |
| xos0063370 | PTP Delay Response correction field has high value (Junk value) when PTP Delay response packets were passing through combo ports. |
| SummitStack | |
| xos0061777 | Standby nodes do not come back up to operational state after they go into failed state. |
| xos0062853 | Communication failure occurs between hosts connected to SummiStack after the stack is broken into two separate stacks, and then merged back to single stack. |
| xos0063490 | CFM stays down after slot reboot on a stack. |
| xos0063788 | The following error appears continuously in backup/standby nodes when node is put in the failed state due to a license/ExtremeXOS mismatch: <pre>&lt;Error:DM.Error&gt; Slot-2: Node State[185] = FAIL (License Mismatch)</pre> |
| Summit X450-G2 Series Switches | |
| xos0060129 | On Summit X450-G2 series switches, 10/100/1000BASE-T SFP+ optics do not link to similar optics when in the SFP/SFP+ ports. They do link or partially link when connected to a regular triple speed copper port. |
| Summit X460 Series Switches | |
| xos0063595 | On Summit X460 series switches, the command to configure stacking ports does not show the native option. |

Table 9: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3-Patch1-2 (continued)

| CR Number | Description |
|---------------------------------------|--|
| Summit X460-G2 Series Switches | |
| xos0062913 | On Summit X460-G2 series switches, copper combo port does not advertise its flow control capabilities to peers. |
| xos0063495 | Policy authentication fails when radius request queue has stale entries. |
| xos0063927 | Error "Deferred L2 notification code out of sync unit 0" repeatedly appears in log. |
| Summit X670 Series Switches | |
| xos0057671 | Link status goes to Ready state on port with 10/100/1000BASE-T optics after multiple reboots. |
| Summit X670-G2 Series Switches | |
| xos0063665 | When Tenant VM connected to MLAG (LACP) ports is turned off, tenant VM FDB is not aged out even after the age out timer expires. The FDB entry persists. |
| xos0063807 | On Summit X670-G2 series switches, egress ACL rule actions do not take effect on ports 64-72 |

Resolved Issues in ExtremeXOS 16.1.3

The following issues were resolved in ExtremeXOS 16.1.3. ExtremeXOS 16.1.3 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, and ExtremeXOS 16.1.2. For information about those fixes, see the release notes for the specific release.

Table 10: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.3

| CR Number | Description |
|----------------|---|
| General | |
| xos0061027 | For SummitStacks, creating or deleting non-default QoS profiles may cause some ports to flap. |
| xos0063274 | VLAN packets are egressing with VMAN ethertype when an egress port is deleted from a VMAN that is also part of a VLAN. |
| xos0063368 | In an MLAG configured switch, FDBs are not installed in hardware after reboot if there are frequent MACMoves between MLAG port and ISC. |

Resolved Issues in ExtremeXOS 16.1.2-Patch1-4

The following issues were resolved in ExtremeXOS 16.1.2-Patch1-4. ExtremeXOS 16.1.2-Patch1-4 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, and ExtremeXOS 16.1.2. For information about those fixes, see the release notes for the specific release.

Table 11: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2-Patch1-4

| CR Number | Description |
|----------------|---|
| General | |
| xos0063359 | The process <code>rtmgr</code> might end unexpectedly after executing <code>disable bgp</code> , and then <code>enable bgp</code> , or after <code>disable port</code> , and then <code>enable port</code> , or after rebooting a switch containing BGP routes. |
| xos0063506 | Traceroute MAC address in CFM domain does not return information about destination switch. |
| xos0063380 | Error message appears after rebooting switch with OSPF configuration: "Error while loading "ospfInterface": ERROR: 0.0.0.0 is not a valid configured neighbor for interface". |
| xos0057269 | SNMP trap <code>extremelpSecurityViolation</code> is sent with incorrect VLAN description. |
| xos0061507 | SNMPget on EXTREME-SOFTWARE-MONITOR table returns value with incorrect OID. |
| xos0061781 | Identity manager entries become stale when clients are moved from one port to another in sub-VLANs. |
| xos0061788 | The process <code>devmgr</code> ends unexpectedly during <code>snmpwalk</code> when continuous EMS logs are sent to the switch console. |
| xos0061797 | Dot1x client moves to authentication failure VLAN if authentication failed due to incorrect supplicant password or framework failure, such as error in VLAN movement, etc.; even if web-based NetLogin is enabled. |
| xos0061855 | Configured OSPF neighbor is not retained after rebooting. |
| xos0062240 | Port that was administratively disabled becomes up after enabling rx pause. |
| xos0062366 | After reboot, DHCP binding entries are not restored via <code>vr-default</code> . |
| xos0062537 | HAL crash occurs when <code>redirect-port-list</code> action contains more than 64 ports. |
| xos0062618 | ELRP forgets the disabled port information if the port is deleted from another VLAN that also has ELRP enabled. As a result, the disabled port stays disabled unless manually enabled. |
| xos0062619 | SSH access-profile using policy does not work with IPv6 addresses. |
| xos0062629 | Clearflow rule does not work properly if there is <code>dot(.)</code> in the ACL counter. |

Table 11: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2-Patch1-4 (continued)

| CR Number | Description |
|--|---|
| xos0062879 | Transceiver information shows same Rx power value for 4x10G partition ports even though some ports are in ready state. |
| xos0063089 | Kernel oops triggered infrequently during continuous addition/deletion of ARP entries for long durations. |
| xos0063108 | NAS Identifier attribute should be sent in RADIUS accounting requests. |
| xos0063120 | Error message "CFP2 modules >= 18 W unsupported" incorrectly appears for Finisar Corp CFP2 LR4 optics. |
| xos0063172 | ACL action "redirect-port-list" does not take effect when another slice has a rule to match all packets with deny action. |
| xos0063248 | NTP MD5 authentication with NTP server is failing. |
| xos0063257 | Saving configuration fails/times-out when VLANs added to a mirror filters are renamed. |
| xos0063271 | Layer 3 packets in non-default virtual routers are slow-path forwarded after disabling MPLS in the peer switch. |
| xos0062494 | Source MAC addresses learned through MVRP packets on a blocked port (STP) cause traffic to be dropped. |
| xos0062701 | HAL timeout occurs while rebooting a stack with STP configuration. |
| xos0062754 | VPLS traffic egresses out with dot1q tag when secondary EtherType is configured. |
| xos0063090 | Netlogin client does not move into authfail VLAN when user is absent from local database. |
| xos0063207 | Error occurs while adding LAG ports as tagged in one VMAN and untagged in another VMAN, even though the VMAN EtherType is primary. |
| xos0062821 | ACL rules installed are not mapped to single virtual group even though ACL action-resolution mode is highest-priority. |
| BlackDiamond 8800 Series Switches | |
| xos0062009 | In BlackDiamond 8800 series switches with XL modules, clearing FDBs when there is a loop causes the FDBs to lose synchronization across slots or switching units. |
| SummitStack | |
| xos0063242 | Stacks configured as DHCP clients do not respond to pinging after failover. |
| xos0062949 | HAL process ends unexpectedly in stack after executing the following commands <pre>debug hal configure stacking pdu-trace mask 0xf debug hal configure stacking pdu-trace capture cap_file</pre> |
| xos0063344 | With MLAG and LAG configurations, when a stack node comes up after a reboot, FDB entries flooded from other slots are programmed on wrong ports internally. |

Table 11: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2-Patch1-4 (continued)

| CR Number | Description |
|---------------------------------------|---|
| xos0063349 | Switch stops responding to SNMP requests if SNMP get for multiple OIDs is continuously initiated. |
| xos0062367 | ACL process ends unexpectedly on repeated refresh of ACL policy with clear-flow action. |
| Summit X440 Series Switches | |
| xos0062621 | On Summit X440-8p switches, the <code>show fan</code> command output displays that the fan is unsupported. |
| Summit X460 Series Switches | |
| xos0063206 | Cannot add L2 entries in hardware due to a full L2 table caused by hash collisions. |
| Summit X670 Series Switches | |
| xos0059514 | On Summit X670V-48x switches, after multiple reboots, 40G ports can remain in ready state, rather than coming to active state, which impacts the traffic passing through those ports. |
| xos0061559 | Enabling OpenFlow on VLANs causes double-wide ACL slice to be used even though it can fit in single-wide slice. |
| xos0062487 | QSFP+ optics are detected as unsupported after rebooting. |
| xos0063137 | Known unicast traffic is not shared between the stacking high-gig trunk ports. |
| Summit X670-G2 Series Switches | |
| xos0061791 | On SummitStacks containing master and standby nodes of different switch platforms, the standby node may go to failed state after a node reboot. |
| xos0063204 | Traffic stops on LAG ports when frequently modifying the sharing group. |

Resolved Issues in ExtremeXOS 16.1.2-Patch1-1

The following issues were resolved in ExtremeXOS 16.1.2-Patch1-1. ExtremeXOS 16.1.2-Patch1-1 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, ExtremeXOS 16.1, and ExtremeXOS 16.1.2. For information about those fixes, see the release notes for the specific release.

Table 12: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2-Patch1-1

| CR Number | Description |
|----------------|---|
| General | |
| xos0059165 | OSPFv3 crash occurs after deleting newly added VLANs. |

Table 12: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2-Patch1-1 (continued)

| CR Number | Description |
|------------|--|
| xos0059560 | After reboot, OSPFv3 fails to select the best path to destination. |
| xos0059569 | OSPFv3 external routes flap and traffic loss occurs when introducing new ABR which has reachability to destination. |
| xos0061505 | After a topology change in the network, BGP routes requiring two levels of recursive lookup are programmed in hardware with incorrect next hops. |
| xos0062017 | DHCP trusted port configuration is lost after disabling, and then re-enabling LAG. |
| xos0062045 | LLDP packets are tunnelled over L2VPn. |
| xos0062290 | Due to ExtremeXOS reflection RSTP BPDU support, upstream bridges believe that they are receiving their own BPDUs (contain the bridge's ID), thus causing multisource events during topology changes, which can cause slow convergence times when Ip is configured (upwards of 30 seconds). |
| xos0062428 | Member ports with a modified speed configuration that is different than the master port should not be allowed in LAG. |
| xos0062441 | The process rtMgr ends unexpectedly when IPv6 static route is deleted. |
| xos0062460 | The <code>show configuration</code> command output shows incorrect ELRP configuration |
| xos0062674 | UPM profile fails to set the variables received from the RADIUS server using VSA 212. |
| xos0062709 | Due to the stronger hash algorithm, if you create accounts in ExtremeXOS 16.1, and then downgrade to versions earlier than ExtremeXOS 16.1, you may encounter problems using the passwords for these accounts. For more information about this issue, visit: http://extr.co/1KfSszY . |
| xos0062719 | Allow use of 3rd-party optics without any additional license. |
| xos0062728 | OSPFv3 best path is not selected after issuing <code>restart ports all</code> in peer switch. |
| xos0062756 | Output of <code>show network-clock gtp current-set</code> command shows incorrect timestamp for "Last GM Change Event". |
| xos0062789 | Disabling learning on LAG ports does not flush FDB entries. |
| xos0062914 | The process mcmgr ends unexpectedly after receiving corrupted IGMPv3 join packets on MLAG ports. |
| xos0062965 | Policy process ends unexpectedly with signal 6 when master node goes down. |
| xos0063071 | Add support for ONEPolicy IP socket classification. |
| xos0054348 | Cannot delete flow names after deleting, and then creating, the flow while the ACL is installed. |
| xos0057538 | OSPFv3 fails to select the best cost external route. |
| xos0057574 | After multiple disable/enable OSPFv3, OSPFv3 routes are not advertised to route manager. |
| xos0057575 | OSPFv3 external routes are not updated in the routing table after link flap events. |
| xos0057583 | In OSPFv3, after adding an ASBR, new route is received, but it is not added to routing table even though it is a best route. |

Table 12: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2-Patch1-1 (continued)

| CR Number | Description |
|--|---|
| xos0057584 | When there are equal cost ASBR routes, OSPFv3 is sharing only one path to route manager. |
| xos0059341 | OSPFv3 external routes are flushed after interface cost is changed if multiple ABRs are present for an area. |
| xos0059446 | OSPFv3 external routes are flushed when ports from OSPFv3 user VR are deleted. |
| xos0060463 | OSPFv3 external routes are flushed after the command <code>restart ports all</code> is executed in area border router. |
| xos0061890 | Zero-Touch Provisioning should not run when default.xsf file is present. |
| xos0061965 | Configuring ESRP member VLANs (VRRP-enabled) produces errors. |
| xos0062380 | Switch rejects incorrect LSP configurations as expected, but this operation still uses LSP indexes in hardware. |
| xos0062504 | You can set a GPTP "peer delay current interval" outside of the correct range of -3 to 17. |
| xos0062508 | When a port is added in a loopback VLAN, OSPFv3 route is not advertised with /128 mask. |
| xos0062824 | Extra ACL slice allocated for QoS even with default configuration. |
| Summit X430 Series Switches | |
| xos0059486 | On Summit X430 series switches, optics are not detected after repeated removal and reinsertion of optics when CPU is busy. |
| SummitStack | |
| xos0062800 | Stack node fails because of license mismatch for 3rd-party optics. |
| Summit X460-G2 Series Switches | |
| xos0062855 | On the Summit X460-G2 series switches, VPLS packets are forwarded with two tags when the service VLAN ports are also members of an untagged VMAN. |
| Summit X670 Series Switches | |
| xos0063052 | Traffic loss occurs on computer connected to Summit X670v-48t switches when the connected switch port is oversubscribed in 100 MB mode. |
| BlackDiamond 8800 Series Switches | |
| xos0061663 | In c- and e-series I/O modules after MLAG port up/down events, it takes long time to learn FDB entries. |
| BlackDiamond X8 Switches | |
| xos0063057 | On BlackDiamond X8 series switches, an IPv6 multicast group, such as FF02::1, is incorrectly treated as a local address. |

Resolved Issues in ExtremeXOS 16.1.2

The following issues were resolved in ExtremeXOS 16.1.2. ExtremeXOS 16.1.2 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-

patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, and ExtremeXOS 16.1. For information about those fixes, see the release notes for the specific release.

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2

| CR Number | Description |
|----------------|---|
| General | |
| xos0057374 | Switch odometer value is reinitialized when Master Switch Fabric Module (MSM) fails to read the value. |
| xos0058669 | DHCPv6 client: After changing the client identifier type, and then restarting the port, old IPv6 addresses are not released, causing the show vlan command to show multiple IPv6 addresses. |
| xos0059942 | SSH connection ends when show commands produce lengthy output. |
| xos0059989 | Configuring non-persistent command using UPM script shows dirty bit(*) in the prompt. |
| xos0060909 | In UPM profiles the variable EVENT.TIME incorrectly has the current time rather than the time when the event was queued/triggered. |
| xos0061085 | Kernel oops occurs while deleting VR with enable BGP export and IPARP proxy configurations. |
| xos0061173 | L2PT packets are dropped when ingress port is configured with software learning. |
| xos0061219 | Parallel-mode-enabled DHCP offer is sent using primary IPv4 address to the client for multiple offers received from server for different IPv4 addresses. |
| xos0061331 | Bootprelay for VRF is not supported. Commands to configure bootprelay should reflect this. |
| xos0061379 | Switch temperature value retrieved using SNMP get operation is incorrect. |
| xos0061445 | After creating and enabling an STPD, the command <code>configure "Default" add ports 1 tagged stpd "s1"</code> adds ports to the Default VLAN, but not with STPD domain, even though the error Command Aborted and no changes were made appears. |
| xos0061565 | The TCL function, "clock scan," generates errors with default time zone configuration. |
| xos0061656 | Nodes remain in the "FDBSync" state due to temp-flooding while rebooting the stack. |
| xos0062018 | For IPv6 routes with mask lengths greater than 64-bits, IPv6 unicast packets destined for the switch CPU can be dropped if another IPv6 route is present with a matching prefix and mask length less than or equal to 64-bits. This issue affects Summit X440, X460-G2, X670-G2, X770 switches, BlackDiamond 8800 G48Te2 I/O modules, and BlackDiamond X8 100G I/O modules. |
| xos0062128 | L3VPN traffic is not forwarded after executing <code>disable port</code> and <code>enable port</code> in MPLS core network. |
| xos0062133 | STP flush event does not happen after ports are quickly disabled, and then enabled. |
| xos0062271 | CLI memory leak occurs when executing show commands with include option through script. |
| xos0062427 | EDP process ends unexpectedly when CDP packets without portId TLV are received. |

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2 (continued)

| CR Number | Description |
|--|--|
| xos0062472 | Source MAC addresses learned through CDP packets received on EAPS-blocked ports cause traffic to be dropped. |
| xos0062710 | On BlackDiamond 8800 or BlackDiamond X8 series switches, with Distributed IP ARP mode on, ECMP routes are sometimes not installed when gateways flap. |
| BlackDiamond X8 Series Switches | |
| xos0058950 | XFP optics appear as unsupported in the output of the <code>show port configuration</code> command. |
| xos0060085 | The command <code>show iproute reserved-entries statistics</code> displays the incorrect maximum value for IPv4 local hosts: "16,384" versus the correct value of "65,533" for BDx8 XL modules. |
| xos0060264 | The output of the <code>show port transceiver info</code> command for optics inserted in 40G/100G ports might be abnormally lengthy if the same command is executed from two different CLI sessions simultaneously. |
| xos0061186 | Bytes counter associated with <code>show port utilization</code> command output displays inaccurate value for 100G ports when utilization exceeds 40%. |
| xos0061639 | Packets ingressing on VLAN-bridged interfaces (Layer 2 VLANs) are not forwarded when the destination MAC address is the same as the switch MAC address, and the switch has at least one Layer 3 interface. |
| xos0061902 | BlackDiamond X8 series switches use VLAN instance as index instead of router interface (rtif) for ARP entries. |
| xos0062306 | Packets get reflected with same tag on port-specific, tag-enabled VLANs after failover. Issue happens only on switches having both port-specific tags and MPLS RSVP-TE configurations. |
| xos0062477 | BlackDiamond X8 series switches' management ports flap and show <code>Detected Tx unit hung</code> error messages. |
| xos0062499 | Multicast packets are dropped in Layer 2 bridged VLANs. |
| BlackDiamond 8800 Series Switches | |
| xos0060891 | While sending continuous join and leave, as well as multicast streams multiple times allowing the streams to age out, kernel error messages appear for add/delete operation with reason <code>Entry not found</code> . |
| xos0061822 | HAL process ends unexpectedly during failover when switches have ACL policies without meter action. |
| xos0061994 | Packets are not forwarded over VPWS after rebooting BlackDiamond 8800 switches. |
| xos0062060 | FDB entries are not programmed in hardware even though hardware resources have sufficient capacity. |
| xos0062535 | Matched packets are dropped for a subset of rules while using <code>redirect-port-list</code> action. |
| Summit Family | |
| xos0059007 | For Summit X670-G2-48x and X770 switches, QSFP+ to SFP+ adapter support is added to work with all optical SFP+ transceivers with the exception of LRM and passive copper direct attach cables. |

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2 (continued)

| CR Number | Description |
|------------------------------------|--|
| xos0062782 | Some existing Summit X430 or X440 configurations requiring Access Lists may stop working after upgrading to ExtremeXOS 15.7.2 or 15.6.3. |
| xos0062113 | The <code>show power</code> command output does not display power usage for PSUs with part numbers starting with "800515". |
| xos0061886 | SNMP master process ends unexpectedly with signal 6 with certain sequence of <code>snmpbulkget</code> and <code>snmpget</code> . |
| SummitStack | |
| xos0062570 | In SummitStacks, executing the command <code>enable sflow ports all</code> enables sFlow inappropriately on stacking ports. |
| xos0057835 | In SummitStacks, clear-flow sampling period is incorrectly calculated. |
| xos0061108 | With authentication mode set to optional, MAC addresses exceeding the allowed number of users are not learned in the However, FDB. Traffic is forwarded correctly by the default policy. |
| xos0061614 | On Summit X450-G2/X440 mixed stacks, executing the command <code>show stacking</code> causes the master to become unresponsive after rebooting backup. |
| xos0061784 | Failover with ONEPolicy enabled with thousands of user sessions authenticated may result in some sessions being tracked incorrectly on the new master, which may result in inconsistent session behavior for those sessions. |
| xos0061799 | Precedence order between policy port rules and policy MAC-based rules is not preserved following a master/backup Failover. |
| xos0061847 | If ONEPolicy is enabled and Summit X670-G2 or X770 series switches are inserted and/or a slot is rebooted, the Summit X670-G2 or X770 fails to properly enable ONEPolicy. An error message similar to the following appears: <pre>mm/dd/yyyy hh:mm:ss.ms <Error:HAL.IPv4ACL.Error> Slot-2: pibAclWrapReserveSlices failed for slot 1 unit 1 with rv=-14</pre> |
| xos0061957 | HAL process ends unexpectedly during failover when switches have ACL policy without meter action. |
| xos0062217 | In SummitStacks with eight nodes and sFlow configuration, Hardware L3 Table full error messages appear when the stacks have a large number of Layer 3 entries. |
| xos0062238 | On a stacked system, configuration of a user-defined CoS value's <code>etsysCos8021dPriority</code> using the MIB can return success when the set actually failed (as seen by a subsequent get). |
| xos0062522 | In Summit stacking switches, standby slots go to failed state when very large log messages are continuously generated in the switch. |
| Summit X430 Series Switches | |
| xos0061864 | FDB process consumes more than 20% utilization when Summit X430 switches are configured with 100+ VLANs. |
| Summit X440 Series Switches | |
| xos0061578 | Configuring NTP on VR-Mgmt is not allowed in virtual router unsupported platforms. |

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2 (continued)

| CR Number | Description |
|---------------------------------------|--|
| Summit X450-G2 Series Switches | |
| xos0061070 | Summit X450-G2 1G switches (X450-G2-xxx-GE4) display incorrect log messages when SFP+ optics are inserted into any of the 1G SFP ports. Incorrect Log message example: <code>The configuration for the SF+_LR optic module is not correct - please configure port 50 for auto-negotiation Off and speed 10000</code> |
| xos0061933 | Summit X450-G2 series switches may overheat. |
| Summit X460 Series Switches | |
| xos0061180 | In Summit X460 stack with mixed alternate and native stacking enabled slots, traffic ingressing one specific slot is not forwarded to other slots. |
| xos0061599 | On Summit X460 series switches, after configuring limit learning, traffic is doubled for the non-blackholed MAC addresses, whereas traffic is blocked correctly for blackholed MAC addresses. |
| Summit X460-G2 Series Switches | |
| xos0061486 | Combo ports have unsupported autonegotiation and half-duplex settings. |
| xos0061814 | When NetSight Policy Manager attempts to enable policy, Policy Manager might report success even if there is a failure. The failure is reported properly by EMS message. |
| xos0061892 | When policy profile has CoS status enabled without pvid-status enabled, disabling policy does not clean up underlying rules; subsequently, enabling policy may fail. |
| xos0062425 | On Summit X460-G2 series switches, the primary port is incorrectly set as 40 when it should be 41. Under certain conditions, this can cause a kernel crash. |
| Summit X670 Series Switches | |
| xos0061167 | Links become active without a connection with tri-speed Base-T SFP installed. |
| Summit X670-G2 Series Switches | |
| xos0061818 | After rebooting Summit X670-G2 stacks, links come up after long delay of ~20 minutes when QoS profile is configured on 500+ VLANs. |
| ACLs | |
| xos0059924 | The output of the command <code>show access-list meter ports</code> displays additional meter name when only one meter is applied using ACL policy. |
| xos0060716 | Need support for new ACL action "redirect-vlan" to redirect matched packets to all ports in specified VLANs. |
| xos0061922 | Dynamic ACLs applied as "any" fail to install in hardware after upgrading ExtremeXOS from any release other than EXOS 15.3. |
| xos0062145 | With QoS configuration, ACL process signal 11 ends unexpectedly after rebooting. |
| BGP | |
| xos0061411 | Route table installs sub-optimal BGP routes (next-hop) to kernel, while the BGP RIB shows different paths when same routes are received from two different peers in local-RIB. |

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2 (continued)

| CR Number | Description |
|----------------------------|--|
| xos0062260 | BGP process ends unexpectedly when local address or password is changed for BGP neighbor, and then you immediately execute a BGP show/configuration command. |
| Chalet | |
| xos0060354 | ExtremeXOS Chalet using IPv6 does not work with HTTPS. |
| xos0062016 | Command line process memory leak occurs when accessing switches with Chalet. |
| EAPS | |
| xos0061038 | Loops occur in EAPS-protected VLANs, after peer reboot, if a VLAN's port is also protected by ELSM. |
| xos0061385 | EAPS process ends unexpectedly after deleting EAPS shared-port configuration. |
| FDP | |
| xos0059655 | Error Unable to delete permanent entry appears while deleting static blackhole FDB entries. |
| Identity Management | |
| xos0061699 | Traffic is dropped when moving idmgr client from one port to another with role-based authentication. |
| IP General | |
| xos0057672 | The process rtmgr ends unexpectedly when disabling GRE tunnels. |
| xos0060796 | Loop created for IGMP query and report packet when port isolation is "ON". |
| xos0061129 | In a multi-peer setup with many routes (over 150K), a few routes from the preferred peer do not become active in the BGP RIB. Disabling, and then re-enabling peer, restores all routes. |
| xos0061198 | Disabling VPN-VRF affects traffic on another VPN-VRF. |
| xos0062705 | Kernel oops can occur after clearing IPMC FDB in a stack. |
| MPLS | |
| xos0061092 | Traffic forwarding on VPLS-serviced VMAN stops after link flap. |
| xos0061662 | When a large number of RSVP-TE LSPs (>250) are signaled to go down because of a link failure, some LSPs do not go down right away. They time out after a few minutes. |
| NetLogin | |
| xos0058808 | Rarely, MAC addresses of authenticated clients learned on NetLogin-enabled ports are not programmed in hardware. |
| xos0060242 | Ports are not removed when unauthenticated UPM profiles are triggered. |
| xos0060449 | NetLogin MAC-based mode does not work as expected with PVLANS. |
| xos0061450 | Logging on to the switch using HTTPS is permitted if NetLogin is enabled, even though HTTPS is disabled. |
| xos0061820 | Dot1x clients move to authentication failure VLAN when web-based NetLogin is enabled globally. |

Table 13: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs) for ExtremeXOS 16.1.2 (continued)

| CR Number | Description |
|--------------|---|
| xos0061837 | When a user is authenticated by NetLogin, but exceeds the allowed per-agent user limits in policy, policy does not signal NetLogin to unauthenticate the user. |
| xos0061868 | With protocol order as MAC dot1x, web-based UPM profile is not executed for the client, which is authenticated as MAC. |
| QoS | |
| xos0060092 | Fetching values using SNMP for "extremePortQosStatsTable" does not work correctly. |
| SNMP | |
| xos0059964 | SNMP poll for MIB dot3StatsDuplexStatus always returns unknown(1) when ports are configured with auto-negotiation on. |
| xos0061945 | SnmpSubagent crash occurs when snmpset executed on the last row in EAPSMbrVlanEntry. |
| VLANs | |
| xos0054039 | IP multicast traffic is not forwarded on PSTAG VLANs when it shares ports with other IGMP snooping-enabled VLANs or other L3 VLANs. |
| xos0061840 | A large VLAN configuration (near capacity for the switch) can stop working after restarting a VLAN client application (for example, <code>restart process eaps</code>) or rebooting due to excessive memory usage. |
| xos0062277 | The command <code>show vlan vlan_list</code> does not show information for dynamic VLANs nor the Default VLAN. Error appears. |

Resolved Issues in ExtremeXOS 16.1

The following issues were resolved in ExtremeXOS 16.1. ExtremeXOS 16.1 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, and ExtremeXOS 15.7.1. For information about those fixes, see the release notes for the specific release.

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1

| CR Number | Description |
|----------------|--|
| General | |
| xos0059913 | After many authentications and aging out of NetLogin clients, some ports are not learning FDB entries. |
| xos0051961 | Unable to block IPv6 traffic from SSH/Telnet/Web interface by access-profile policy. |
| xos0055108 | The bound IP address is not being reflected in the command <code>show vlan</code> . |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|------------------------|--|
| xos0055358 | Device manager is reporting incorrect slot type after clearing a slot while that slot is disabled. |
| xos0055399 | After multiple failures in a main ring, the Ring Protection Link (RPL) of the sub-ring stays in the blocked state in an RPL neighbor node of the sub-ring causing traffic to fail between the nodes in the sub-ring. |
| xos0056987 | With ELSM enabled on port p1 on two switches, with 10 seconds as the ELSM hello interval, and trap receiver configured on the both switches (or managed in Ridgeline, which configures the trap receiver by itself), and then disable/enable port p1 on both switches. Switches send two linkup traps. |
| xos0057211 | Traffic gets forwarded for blackholed MAC address when limit learning is enabled. |
| xos0057226, xos0057225 | Creating an account with a 32-character long name, and then opening an SSH session from that user, causes the switch to crash. |
| xos0057254 | For every reauthorization of NMS clients the policies attached are unbound and bound back again in the peer (not in the authenticator). |
| xos0057297 | When MD5 password is different for both broadcast server and client, the NTP session is established anyway, and no warning is indicated in "leap indicator". |
| xos0057391 | Configured restrict-list entry is shown as user, and active peer interface is removed from NTP peer when its IP address is configured in deny option. |
| xos0057392 | NTP session still shows synchronized after disabling port in NTP client interface. |
| xos0057462 | SSL certificate generated in switch uses a weak hashing algorithm, MD5, which is known to be vulnerable to collision attacks. |
| xos0057463 | The minimum key length supported when configuring an SSL certificate is 64 bits, which is considered medium strength and could be exploited by an attacker on the same physical network. |
| xos0057517 | With global NTP server configured and an NTP session established between server and broadcast client, after disabling NTP, the peer loses its active peer status (*), but for the broadcast client, status (*) is not removed. |
| xos0057547 | The output of the <code>show configuration nettools</code> command does not show IPv6 smart relay configuration for VLANs in user-created virtual routers. |
| xos0057570 | Executing the command <code>configure bootprelay ipv6 prefix-delegation snooping add</code> , the ipv6-prefix valid timer alone is updated and gateway is not updated for the same prefix, which is in the prefix delegation snooping table. |
| xos0057590 | Intermittently, 50% traffic loss is seen after one MLAG peer is rebooted. |
| xos0057599 | After giving gateway as the same interface address on VLAN in <code>config bootprelay ipv6 prefix-del snooping add</code> command, error does not occur and the same route is added in snooping table and rtmgr table. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|------------|--|
| xos0057746 | Dynamic routes learned using bootp for indirect routes are not deleted after disabling the baseline protocol (Ripng/Static route) for the gateway. |
| xos0057849 | Client session using ciphers aes128-ctr, aes192-ctr, aes256-ctr is not working. |
| xos0057976 | DHCP bootprelay option 82 with aggregation port shows incorrect circuit ID. |
| xos0058324 | Add port number to output of all <code>show port</code> commands along with existing display string. |
| xos0058325 | Add multicast tx counter to <code>show port stat</code> command. |
| xos0058326 | Add a flag that indicates MLAG membership to the output of the <code>show port info</code> command. |
| xos0058327 | In <code>show port port</code> info detail command, order the VLAN listing alphabetically. |
| xos0058337 | Generate Syslog message or trigger an SNMP trap when the log memory buffer is filled up to 90% of configured log lines. |
| xos0058349 | Enabling DHCP on any VLAN with IP address already configured does not produce error. |
| xos0058384 | XNV Dynamic VLAN: With VM-tracking enabled on MLAG peer ports, dynamic VLANs are not created for tagged and untagged traffic. |
| xos0058393 | The TCL command <code>clock format</code> is not available for CLI scripting and attempting to use it produces an error. |
| xos0058668 | After rebooting DHCPv6, client stuck up in rebooting state. |
| xos0058719 | Dynamic routes added by prefix delegation snooping remain active even after disabling and unplugging the client connected to the port. |
| xos0059173 | Traffic loss occurs since the destination IP address is missing in kernel after port fail. |
| xos0059222 | SFLOW-sampled packets are flooded out of VLANs when these same packets are software learned. |
| xos0059243 | The process <code>exsh</code> ends unexpectedly after executing a show command with a port list followed by invalid letters (for example, <code>show port 1:1, 1:2ab</code>), and then pressing [Tab] . |
| xos0059247 | ARP entries incorrectly point to ISC port after MLAG peer is rebooted. |
| xos0059315 | Unable to add static routes with tunnel address as the next-hop. |
| xos0059524 | Link status is incorrect when auto-polarity setting is off. |
| xos0059558 | XNV Dynamic VLAN: Running the command <code>unconfigure vm-tracking local-vm mac-address 00:00:00:00:00:01 vlan-tag</code> , FDB is not removed from dynamic VLAN and the command <code>show vm-tracking</code> displays no entries. |
| xos0059581 | Rtmgr process ends unexpectedly when OSPF external routes are deleted from the route table. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|--|--|
| xos0059597 | Static ARP entries are created as blackhole in HAL IPV4adj, thus the next hop information for the routes with those adjacency as gateway are not programmed properly in I3 defip table. |
| xos0059661 | Running extended diagnostics on backup MSM (Master Switch Fabric Module) can, under certain rare conditions, cause the <code>cfmgr</code> process to end unexpectedly on the master MSM. |
| xos0059851 | When a DHCP client receives an IP address that conflicts with a static IP address configured in a VLAN, the static IP address is removed from the VLAN and the DHCP client stops. |
| xos0059952 | XNV Dynamic VLAN: After disabling dynamic VLAN on vmt-enabled port, the "v" flag is not removed from the output of the <code>show fdb</code> command. |
| xos0059983 | RADIUS shared secret password obfuscation is too weak. |
| xos0060088 | Kernel oops triggered rarely during continuous addition/deletion of ARP entries for long duration in presence of high CPU utilization. |
| xos0060100 | Kernel oops occurs due to memory corruption caused by slow-path forwarded traffic. |
| xos0060119 | Changing the primary TACACS server configuration locks out TACACS-authenticated users. |
| xos0060228 | HAL process ends unexpectedly in rare circumstances while rebooting switch with default speed configuration on 10G ports |
| xos0060228 | HAL process ends unexpectedly in rare circumstances while rebooting switch with default speed configuration on 10G ports. |
| xos0060478 | There is no command to configure the idle timeout of SFTP sessions. The idle timeout is fixed at 60 minutes. |
| xos0060713 | Executing the command <code>clear cpu-monitoring</code> sets the dirty bit(*). |
| xos0060825 | Double-tagged CFM frames are dropped by kernel in VMAN environment. |
| xos0060977 | MPLS traffic is egressing switch with incorrect label causing packet loss. |
| xos0061178 | Dynamic ACL for gratuitous ARP violation on LAG member ports are incorrectly getting installed on LAG master ports. |
| xos0061517 | LACP Adjacency failed while forwarding the PDU with I2pt profile over L2VPN tunnel when MPLS PHP enabled. |
| BlackDiamond 8800 Series Switches | |
| xos0061338 | PSTAG packets are fail over fabric link. |
| xos0059736 | Process <code>rtmgr</code> ends unexpectedly with signal 11 after executing the command <code>disable bgp</code> . |
| xos0059648 | Static ARP entries are not properly synced with new Master Switch Fabric Module after failover. |
| xos0059009 | After deleting ingress and egress ports from one slot and adding them to another slot, MPLS LSP routes do not get installed in hardware. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|--|--|
| xos0058972 | Layer 2 packets are not forwarded with non-multicast destination MAC addresses for static multiport FDB entries. |
| xos0058895 | With a two-tier MLAG and PVLAN configured, process mcmgr ends unexpectedly with signal 10 and 11 after disabling, and then re-enabling MLAG ports, and disabling remote MLAG peer ISC ports. |
| xos0058776 | Process MPLS crashed with signal 11 after disabling MPLS with broadcast traffic on pseudowire. |
| xos0058394 | BGP routes are missing in XL-series modules after disabling, and then enabling, BGP with 20,000 routes. |
| xos0057624 | Traffic loss occurs on PVLAN after restarting VRRP process in VRRP Master switch. |
| xos0055337 | Issuing the command <code>clear slot</code> on a disabled slot does not clear the pre-existing port numbers associated with the previous module type in that slot. I/O module stays in VLAN sync. |
| xos0050664 | Process DCBGP ends unexpectedly with signal 11 while rebooting neighbor switches in a scaled, multi-homed setup. |
| xos0050732 | On BlackDiamond 8800 series switches, the process DCBGP ends unexpectedly with signal 6/11 after rebooting the switch. |
| xos0054970 | BlackDiamond 8800-xl cards and Summit X480 series switches should not allow Layer 2 Protocol Tunneling and Filtering to be configured over VPLS/VPWS. |
| BlackDiamond X8 Series Switches | |
| xos0060180 | Data traffic over L3VPN: From Core to Access on 100G-XL modules traffic is dropped completely, This issue does not occur while using ports from BDXA-10G48T card. |
| xos0060073 | On the BDXB-100G4X-XL module, known unicast traffic is flooded across all PStag ports of VLAN, instead of going out using one port even though MAC address was learned. |
| xos0059953 | With DHCP client and server configuration, process Nettools ends unexpectedly with signal 6 after running failover. |
| xos0059927 | For the 100G4X and 100G4X-XL modules, tagged data traffic is not going to untagged service VMAN after deleting tagged service VMAN from the VPLS instance, and then creating new untagged service VMAN and attaching it to the VPLS. |
| xos0059710 | Ports on the 100G4X I/O module are reported as having their links down. |
| xos0059534 | Running two failovers leads to traffic drop for some unicast streams. |
| xos0059343 | The process snmpMaster might end unexpectedly during upgrade from ExtremeXOS 15.3 to 15.5 for some SNMP community names. |
| xos0059156 | VRRP control packets are dropped due to congestion in tx queue under scaled environments. |
| xos0056300 | Traffic does not switch back to primary port when smart redundancy is enabled. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|------------|---|
| xos0056773 | Resetting a connected I/O fabric module, produces the following error messages: <pre>04/07/2014 10:09:10.07 <Erro:Kern.Card.Error> Slot-7: pca9506_read, not enough buffer provided (need 5 bytes, have 1 bytes) 04/07/2014 10:09:10.57 <Erro:Kern.Card.Error> Slot-7: i2c-1: bus busy, addr=0x23 rw=1 cmd=0x0 size=2 retry=0 04/07/2014 10:09:10.57 <Erro:Kern.Card.Error> Slot-7: Failed to read data from INPUT_PORT_0_COMMAND (-145)</pre> |
| xos0057560 | Accessing ScreenPlay while running a script can cause the thttpd process to end unexpectedly with the following error: <pre><Erro:DM.Error> MM-A: Process thttpd Failed MM-A rebooted</pre> |
| xos0057827 | After master switch fabric module. (MSM) failover, L2 MC traffic destined to IGMP receivers is not forwarded from ports that are the IGMP member ports sending joins. |
| xos0058375 | ACLs to match VLAN-ID, CVID parameters do not work for slow path forwarded packets. |
| xos0058413 | BDXB-100G4X module reboots with kernel oops with scaled route traffic above 12K IPv4 routes and above 32 IPv4 local hosts. |
| xos0058553 | On BlackDiamond X8 series switches, in large topology (~2,000 pseudowires and lots of LSPs and ILM), disabling all VPLS sessions on switch causes SMAC learning to take longer than usual if the traffic comes from a different pseudowire. Traffic loss occurs during this interval. |
| xos0058615 | ESVT: Loopback port fails to indicate "Link State" of "L". The loopback port is physically looped back. However, if user traffic is in the VLAN, then potentially unwanted traffic can get looped causing issues. |
| xos0058640 | Traffic forwarding inconsistencies occur after restarting MPLS process. |
| xos0058793 | BDX8 XL modules flood decapsulated VPLS traffic, causing bandwidth problems. |
| xos0058814 | Traffic fails to transmit over pseudowire with attachment circuit VMAN carrying CEP-CVID. |
| xos0058842 | Tagged VLAN traffic is dropped at ingress of 100G4X-XL module port, when the same port is added to another CEP-VMAN VPLS. This issue does not occur on non-100G4X-XL modules/ports. |
| xos0058870 | MPLS encapsulation and decapsulation fail when network ports are part of 100G slot. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|-------------------------------|--|
| xos0058979 | With distributed ARP mode on, L3 traffic is slow-path forwarded when next hop entry (IPv4 adjacency) is stored in IPv4 LPM and IPv4 LPM is in external TCAM. If IPv4 adjacency is stored in IPv4 L3 hash table (that is: full LPM or iproute reserved-entries set to maximum), the same traffic is forwarded correctly in hardware. Traffic is also forwarded correctly in hardware if IPv4 LPM table is internal (TCAM mode set to "none"). |
| xos0059104 | ACL policies are not installed in hardware after management module failover. |
| Summit Family Switches | |
| xos0056230 | SNMP query on extremeMemoryMonitorsystemTable does not show backup information, if slot2 is master and slot1 is backup. |
| xos0057106 | When mirroring is configured to be triggered through clearflow, mirroring does not work and produces the following error: 06/26/2014 17:14:16.40 <Warn:HAL.IPv4ACL.Warning> : Could not enable mirroring for ACL rule since mirror acl-rule-1 is not active. |
| xos0058241 | XNV Dynamic VLAN: After disabling, and then enabling, vm-tracking, the <code>show vm-tracking</code> command displays the incorrect number of network VMs authenticated with NMS as the authentication method. |
| xos0058253 | Label-switch paths (LSP) does not work when iproute sharing max-gateways is equal to "4". |
| xos0058841 | NetLogin dot1x IPARP entry is not shown after the iparp detection is turned off and then on. |
| xos0059030 | ARP entries incorrectly point to ISC port after MLAG peer is rebooted. |
| xos0059248 | Identity management: Identity management ends unexpectedly with signal 11 repeatedly when NetLogin MAC addresses are scaled to 512-3,000. |
| xos0059345 | Summit X460-G2-1G switches do not produce warning log messages (not compatible) when 100LX10 optics are inserted in the highest numbered 1G ports. |
| xos0059406 | AVB ACLs are not getting unbound when AVB ports transition from non-boundary to boundary ports. |
| xos0059950 | In Summit series switches, you cannot download bootROM images from memory card. |
| SummitStack | |
| xos0060362 | VPWS/VPLS traffic stops after stack failover, and then failback, when access and core ports are from different slots. |
| xos0060142 | When SummitStack master and backup slots experience prolonged loss of stacking communication (dual master issue), the backup becomes master and later fails due to HAL process ending unexpectedly. |
| xos0060115 | Expected <Info:Kern.Info> log messages are not shown for log filter events. |
| xos0059462 | Timezone configuration is not applied to standby nodes after stack reboot. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|------------------------------------|--|
| xos0059049 | SNMP: Traversing the EntityPhysicalTable takes around about four minutes. |
| xos0058851 | IPFIX egress interface field reports an incorrect value for a SummitStack. |
| xos0050066 | Process ipSecurity PID 1588 signal 11 ends unexpectedly with cleanup session in IPAdresecurity. |
| xos0056075 | After issuing the command <code>restart process ospf-5</code> on a SummitStack, H-VPLS (spoke) nodes fail to encapsulate packets to VPLS pseudowires. Traffic is restored after about 15 minutes. |
| xos0057438 | Memory depletion occurs in Backup/Standby nodes of SummitStack with highly scaled IPFIX flow records. |
| xos0057767 | Static FDB associated with VPLS service VLAN is not programmed in hardware after reboot when "disable learning" is configured. |
| xos0058133 | SummitStacks use slot MAC address in DHCP packets when enabling DHCP on Management VLAN. |
| xos0058145 | Traffic is dropped when enabling MAC-Locking on VMAN CEP ports that are part of the stack backup node. Issue does not occur when ports are on master node in stack. |
| xos0058809 | Process "vlan" does not respond after removing and unconfiguring slot in a stack. |
| Summit X430 Series Switches | |
| xos0057028 | In Summit X430 series switches, kernel gets stuck for few seconds while installing boot v1.0.1.5 from ExtremeXOS. |
| Summit X440 Series Switches | |
| xos0059571 | When creating many VLANs on adjoining devices with MVRP enabled, switch has many dynamic VLANs created on it, causing the switch to reset. |
| xos0059500 | On Summit X440 series switches with more than 1,500 IP ARP entries (exceeding supported hardware limit of ~400), and with ARP entries changing MAC address, some entries are not aged out of hardware. This can cause a mismatch between software and hardware when ARP is relearned with a different MAC address. |
| xos0050402 | The command <code>enable inline-power legacy</code> does not power up pre-standard PoE devices, such as Cisco phone 7940/7960 that do not work with IEEE 802.3af standard detection and legacy capacitive detection. The <code>enable inline-power legacy</code> command now powers up legacy PoE devices that rely on the capacitive detection instead. |
| xos0054084 | In Summit X440 stacks with about 500 identities, restarting the idmgr process, causes the IdMgr process to end unexpectedly with signal 11. |
| xos0058547 | In Summit X440-24t switches, the maximum hotspot temperature should be changed to 70°C. |
| xos0058889 | The output of the <code>show fans</code> command always indicates no fan installed (Empty). |
| xos0059380 | Executing the command <code>!3 !3table hash VRF=0 IP=IPaddr</code> in BCM shell produces a Kernel oops. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|---------------------------------------|--|
| Summit X460 Series Switches | |
| xos0056342 | Misleading power supply unit (PSU) traps are sent when PSUs are inserted or powered on/off. |
| xos0058053 | Summit Stack run failover takes longer than usual time to boot the backup node. |
| xos0058250 | Installing an ExtremeXOS image on a SummitStack after copying the image to the switch using SCP fails when there is no backup module assigned to the stack. |
| xos0059131 | Debounce timer is not getting configured if stack ports reside in different units. Also, pre-emphasis configuration should be rejected in alternate stacking mode. |
| xos0059671 | On Summit X460 series switches with 750 W power supplies installed, log messages Power usage data unknown appear. |
| xos0060057 | Root port only sends STP agreement BPDU for the CIST and first two MSTIs. |
| xos0060517 | When the service VLAN and L2 VMAN (untagged port) is configured on same port, after deleting port from VMAN, service VLAN's traffic is affected. |
| Summit X460-G2 Series Switches | |
| xos0060736 | gPTP propagation delay is not calculated correctly and ports become AVB incapable. |
| xos0061150 | In X460-G2, CLI allows half-duplex config for ports even it is not supported. |
| xos0059827 | During Summit stack failover, Kernel oops appears infrequently, caused by corruption in VR ID of resolved ARPs. |
| xos0059763 | In Summit X460G2-24x-10G switches, the first VIM-2x(10G port 33) sometimes remains in ready state after reboot. This problem does not occur with the second VIM-2x port (34). |
| xos0059652 | Combo copper ports do not work in Summit X460G2-24x-10G4 switches until SFP is also connected. |
| xos0055189 | In Summit X460-G2 stacks, the command show power fails to display power usage and produces the error Failed reading Slot-B power on time during slot reboot. |
| xos0059395 | Summit X460-G2 series switches do not produce warning log message (not compatible) when 100LX10 optics are inserted in the highest numbered 1G ports. |
| xos0055999 | For Summit X670-G2 and X460-G2 series switches, currently the ENTITY-MIB:entPhysicalsFRU for fan and power modules shows the values as FALSE. However since these fan modules and power modules are replaceable units, this should be shown as TRUE. |
| xos0058438 | With MLAG and VRRP IPv6 configuration running for two days, Watchdog timer error, watchdog reboot, and memory depletion occur. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|---------------------------------------|---|
| xos0058517 | The command <code>enable vman cep egress filtering</code> incorrectly drops VLAN traffic if a port is part of both a VMAN and a VLAN. |
| xos0058896 | Running slow path traffic for long durations causes Summit X460-G2 stack instability. |
| xos0059240 | High CPU utilization and MAC learn thrashing occurring on fabric links forming a X460-G2 stack. |
| xos0059577 | On Summit X460G2 series switches, can't install ExtremeXOS SSH XMOD image. |
| Summit X480 Series Switches | |
| xos0059471 | STPID-ELSM: Traffic is not forwarding to STPID "FORWARDING" port in stack/chassis. |
| Summit X670 Series Switches | |
| xos0059128 | In Summit X670 series switch, all LEDs are blinking at a faster rate. |
| Summit X670-G2 Series Switches | |
| xos0060948 | Kernel crashes on Summit X670G2-48x-4q switches when stack ports are configured for v80 stacking. |
| xos0059445 | Link flaps occur when stacks are firmed with 3 m/5 m QSFP+ passive copper cables. |
| xos0058867 | In Summit X670G2-48x-4q stack, system crash occurs with <code>Process ipSecurity pid 1644 died with signal 11</code> error after deleting VLAN with aged timer for ARP entry configuration from IPAddrSecurity. |
| Summit X770 Series Switches | |
| xos0059573 | Factory installed image incorrectly references "x450" in the image name. |
| xos0058536 | In a SummitStack, with X770 and X670-G2 series switches, after a run failover on master Summit X770, backup Summit X770, does not sync. The reverse works fine, and the Summit X670-G2 comes to synced state. |
| xos0055746 | Stacking port link flap occurs on Summit X770 series switches when using 3-meter QSFP+ cables. |
| xos0053310 | On Summit X770 series switches, partitioned 10G ports 77,78,81,82 do not appear in the output of the <code>edp ports all</code> command. Due to this, bi-directional traffic across these ports is not working even though the port is up with speed 10G. |
| xos0055686 | VPLS: IGMP hello packets received on pseudowire are not forwarded to service ports on Summit X770 series switches. |
| E4G-200 Cell Site Routers | |
| xos0060715 | SNMP get does not return the product name for E4G-200-12x. |
| xos0058951 | Executing <code>show ces ces1 details</code> command when CES state is "signaled" causes process <code>ces pid</code> to end unexpectedly with signal 6. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|----------------------------|---|
| xos0058239 | In E4G-200 cell site routers, power supply status displays incorrect value in the output of the show power command. |
| ACL | |
| xos0054720 | With network-zone configured, the command <code>show access-list port 1:1 detail</code> reverses the IP address match criteria. |
| xos0058810 | Both <code>show access-list usage acl-slice</code> and <code>debug hal show device acl-slice slot slot unitunit</code> commands show the same output. |
| xos0059330 | With dual master switch fabric module (MSM) installed, clear-flow ACL intermittently fails. |
| BGP | |
| xos0060670 | Switch becomes unresponsive and resets after issuing command <code>show bgp neighbor</code> . |
| xos0061310 | DCBGP process ends unexpectedly after disabling BGP peer switches after the switch receives VPNv4 routes with the same next hop from two different L3VPN peers. |
| ERPS | |
| xos0059320 | CCM is dropped for "Hardware Down MEPs" when they are received on ports that are blocked by ERPS. |
| ERPS | |
| xos0059146 | With port-specific tags configured, source MAC addresses are removed and re-learned for all incoming ARP packets causing flooded traffic a for short time interval. |
| IP General | |
| xos0050794 | DCBGP process ends unexpectedly with signal 11 when rebooting or issuing the command <code>[enable disable] bgp on</code> neighboring switches. |
| xos0058418 | With inter VR routing configuration, traffic flows at line rate. After disabling the egress port, kernel panic occurs. |
| xos0058432 | In Summit X670-G2 and X770 series switches for internal-table configuration, more L2 IPv4Mcast cache entries are limited to 4,000 instead of 8,000. |
| Identity Management | |
| xos0058112 | The output of the command <code>show identity-management</code> entries displays the ID name and domain name as null. |
| xos0058481 | ACL error occurs when identity management detections are turned on and off. |
| MLAG | |
| xos0056340 | Unknown Layer 2 traffic from Isolated subscriber VLANs are forwarded to the remote MLAG ports, even though local MLAG ports are up. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|-----------------|---|
| xos0058153 | MLAG state flaps continuously when MLAG peers have different versions of ExtremeXOS. |
| xos0060693 | FDB entries in MLAG peers are learned in the incorrect VMAN if the MLAG port is untagged in one VMAN and has CEP CVID configuration in another VMAN. |
| MPLS | |
| xos0060346 | Traffic is not forwarded to secondary path after failover when LSP is added to VPLS. |
| xos0059733 | LSP load sharing does not occur on Summit X460-G2 and BlackDiamond X8-100G4X switches. |
| xos0059729 | Packet duplication occurs after upgrading one of the provider edge (PE) switches in VPLS tunnel. |
| xos0059266 | VPLS: Traffic is dropped after RSVP LSP failover. Failover from primary path to secondary path (same LSP) works. |
| xos0058948 | Changing service VLAN tag results in FDB not being learned. Learning resumes when you revert the VLAN tag. This issue is on both BlackDiamond X8 and Summit X670G2 stacks. |
| xos0058785 | Changing LSR-id of P node causes data traffic drop at ingress of a VPLS provider edge (PE) switch. |
| xos0058502 | Traffic for existing VPLS instance is affected, when a port is removed from untagged VMAN of different VPLS instance. |
| xos0059733 | LSP load sharing does not occur on Summit X460-G2 and BlackDiamond X8-100G4X switches. |
| xos0058460 | Unable to configure multiple LSP from core node to spoke node in Summit X460-G2 and X670-G2 series switches. |
| xos0058395 | STP is not working properly when added to VPLS. |
| NetLogin | |
| xos0061069 | In Netlogin ISP mode, client MAC addresses configured as static FDBs are removed after reboot. |
| xos0060633 | Possible memory leak while cleaning the authVLANlist in NetLogin during unauthorization. |
| xos0059434 | Unable to re-initialize MAC-authenticated client using extremeMacAuthClientInitialize OID. |
| xos0058929 | The number of reported NetLogin authenticated clients reduces to zero even though clients are in fact authenticated. This happens only when the protocol-order is configured as dot1x > mac > web and some re-authentication-timer is configured. |
| xos0058610 | On disabling NetLogin dot1x, several clients' authenticated value is not cleared fully and the command <code>show netlogin dot1x</code> shows the number of clients authenticated as "6" instead of "0". |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|-----------------|---|
| xos0059557 | Executing the command <code>configure netlogin add mac-list 00:00:00:11:11:11 000000111111 ports 1</code> produces the following error: ERROR: Invalid mask length supplied, must be between 1 and 48. Error was thrown and the mac list was not created |
| xos0059592 | NetLogin authentication protocol order resets to default after save, reboot, and then restart. |
| xos0059593 | The <code>show netlogin mac</code> command output no longer shows the default mac mode configuration. |
| xos0058054 | Disabling NetLogin mac,dot1x,web-based is not disabling NetLogin mac,dot1x,web-based. |
| xos0053634 | MAC-lockdown-timeout on user ports does not work as expected if Netlogin is enabled on those ports. |
| OpenFlow | |
| xos0061479 | OpenFlow process signal 6 ends unexpectedly on Summit X460 stacks. |
| xos0060009 | Some flows cannot be installed when OpenFlow FDB table is turned on. |
| xos0059776 | When MS Lync application (MSLync 2.0 solution) attempts to install a LYNC OpenFlow flow with action set queue and set field IP DSCP with Hybrid Mode enabled, a OFPBAC_BAD_QUEUE error occurs. The flow is not installed. The flow is properly installed after rebooting the switch. |
| xos0059621 | Lync should be able to install flows for hybrid VLANs without knowing the VID for those VLANs. |
| xos0059604 | In a stacked environment the Datapath ID changes with every stack reboot. This may cause OneController to go out of sync as the NodeID changes (stale nodes and flows in OneController). |
| xos0059288 | Cannot install multiple XMOD files in BlackDiamond 8800 and BlackDiamond X8 switches due to lack of space. |
| xos0059160 | Sending a flow install with the vlan-id-present field set to false from the OneController causes the flow to fail to install with an ExtremeXOS bad port error message. The flow installs properly by setting the field to true. |
| xos0059139 | OneController uses <code>set_field:0->vlan_vid</code> for strip VLAN action causing an ExtremeXOS error. |
| xos0058978 | Flows with both output port and controller actions cannot be added. |
| xos0058117 | After enabling, and then disabling, OpenFlow globally and on a VLAN with standard mode, you cannot enable learning on the VLAN anymore. |
| Optics | |
| xos0057346 | Link flaps occur when optics such as 10/100/1000 Base-T or 100FX is inserted into Summit X460-G2 and X480 series switches. |
| xos0059579 | SFP+ ports do not link up with active optical breakout cable. Cable is identified as not supported and treated as a 3rd-party cable. |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|---------------|--|
| OSPF | |
| xos0059574 | OSPF packets larger than 8,192 are dropped even with jumbo frame enabled. |
| xos0056243 | OSPF process ends unexpectedly during frequent route re-calculation caused by switch reboot or master switch fabric module. (MSM) failover or BFD flap events. |
| xos0058221 | Rarely, OSPFv3 process ends unexpectedly with signal 11 when link flaps occur. |
| xos0059305 | OSPF consumes a large amount of memory when a large number of Link State Acknowledgment packets are queued up for transmission. |
| PoE | |
| xos0058994 | POE is not delivering power to several model phones when legacy mode is enabled. |
| Python | |
| xos0060600 | EPM process signal 5 ends unexpectedly when creating a new process with a dash in the name. |
| RIP | |
| xos0058683 | RIP packets are dropped when another VLAN has a secondary IP address configured. |
| SNMP | |
| xos0060103 | SNMP walk does not return all VLANs under extremePortVlanStatsTable. |
| xos0058130 | Q-Bridge MIB: On doing SNMP walk for dot1qVlanStaticEgressPorts and dot1qVlanStaticUntaggedPorts for ports in different slots, invalid port number values are returned. |
| xos0058110 | Q-BRIDGE MIB: dot1qPortGvrpStatus and dot1qGvrpStatus should return value as "2" with MVRP disabled globally and at the per port level, but it returns value as "0" which is incorrect as per RFC. |
| TWAMP | |
| xos0060243 | The process rtmgr ends unexpectedly while changing IP address on VLAN when switch is booted with configuration such that an IP route has local or loopback interface address as gateway. |
| xos0060780 | With VRRP enabled, local VLAN's direct route is not installed in hardware after reconfiguring the VLAN's IP address. |
| VLAN | |
| xos0058698 | Traffic stops egressing after adding the egressing port to some other VLAN that is having PSTag configuration with other ports. |
| xos0057435 | Packets are dropped when learning is disabled in a VLAN when its associated ports are configured with limit learning in another VLAN. |
| VRRP | |

Table 14: Resolved Issues, Platform-Specific and Feature Change Requests (CRs) for ExtremeXOS 16.1 (continued)

| CR Number | Description |
|------------|---|
| xos0060794 | VRRP advertisement interval configuration got changed after upgrading EXOS from 12.6 to 15.4 or above releases. Issue happens only when interval is configured in milliseconds. |
| xos0053821 | IPv6 neighbor advertisements for VRRP virtual IP address uses virtual MAC address as source MAC address instead of switch MAC address. |
| xos0061222 | Gratuitous ARP packets for VRRP virtual IP addresses have ARP sender addresses as physical MAC addresses, instead of VRRP virtual MAC addresses. |
| XML | |
| xos0058020 | When both the HTTP and HTTPS are enabled, only HTTPS works, |

4 ExtremeXOS Document Corrections

ACL Policy Redirect
ACL Ports Limits
Configure Sys-Recovery-level Slot Command Platform Availability
Configuring DHCP Binding
ELRP and QoS
ERPS/EAPS Failover Packet Loss in Stacking
Hardware Table Hash Algorithm
L2VPN Sharing Commands
Link Layer Discovery Protocol (LLDP)
Match Conditions Supported for OSPF Import Policy
MIB
MLAG PIM-SM *,G Forwarding Limitation
NetLogin Limitation
NetLogin Local Authentication
PoE Power Delivery
PVLAN + EAPS and PVLAN + STP Recommendations
VRRP Guidelines

This chapter lists corrections to the *ExtremeXOS User Guide* and *ExtremeXOS Command Reference Guide* for

ACL Policy Redirect

In the *ExtremeXOS User Guide*, under chapter **ACL > Policy-Based Routing > Layer 2 Policy-Based Redirect**

xos0062802

The following statement should be removed:

“Using the “redirect-port” action overrides Layer 2 echo kill; the result is that a packet can be made to egress the ingress port at Layer 2.”

ACL Ports Limits

In the *ExtremeXOS User Guide*, under **Comments and Descriptions in ACL Policy Files > Action Modifiers**

xos0062545

The following content:

“redirect-port-list: port_list—Supports multiple redirect ports as arguments. When used in an ACL, matching packets are now redirected to multiple ports as specified in the ACL while overriding the default forwarding decision. (Summit X440, X460, X480, X670, X770, E4G-200, E4G-400, BlackDiamond 8K - 8900-G96T-c, 8900-10G24X-c, 8900-G48T-xl, 8900-G48X-xl, 8900-10G8X-xl, 8900-40G6X-xm, BlackDiamond X8.)”

Needs to be changed to:

“redirect-port-list : port_list—Supports multiple redirect ports as arguments. When used in an ACL, matching packets are now redirected to multiple ports as specified in the ACL while overriding the default forwarding decision. Maximum number of ports that can be mentioned in this list is 64. (Summit X440, X460, X480, X670, X770, E4G-200, E4G-400, BlackDiamond 8K - 8900-G96T-c, 8900-10G24X-c, 8900- G48T-xl, 8900-G48X-xl, 8900-10G8X-xl, 8900-40G6X-xm, BlackDiamond X8.)”

Configure Sys-Recovery-level Slot Command Platform Availability

ExtremeXOS Command Reference Guide for the configure sys-recovery-level slot command.

xos0061985

Under command configure sys-recovery-level slot the existing text:

Platform Availability

This command is available only on modular switches.”

Should change to:

Platform Availability

This command is available on modular and stacking switches.”

Configuring DHCP Binding

In the *ExtremeXOS User Guide*, in **Security > Configuring DHCP Binding**.

xos0064031

The following note should appear:

“When configuring static DHCP binding entries, DHCP binding restoration should be configured.”

ELRP and QoS

ExtremeXOS User Guide, ELRP section.

xos0062230

Add the following note:



Note

ELRP uses QP8 to send/receive control packets over the network to detect loops, so do not use QP8 in the network for any user traffic.

ERPS/EAPS Failover Packet Loss in Stacking

In the *ExtremeXOS User Guide*, in **ERPS > ERPS Feature Limitations**, and **EAPS > monVlanIdStatsTable**.

xos0063703

The following information should appear:

“Using a SummitStack port as the ERPS or EAPS protection port can add approximately 1 second of packet drop time during a failover.”

Hardware Table Hash Algorithm

In the *ExtremeXOS User Guide*, under **Troubleshooting > Modifying the Hardware Table Hash Algorithm > BlackDiamond X8 Switches** section.

xos0061737

The following text:

“Use the following command:

```
disable elrp-client [default | source-port | packet {algorithm [crc | xor] |
dynamic-mode [spray | eligibility | none]}] {slot slot-number}”
```

Should change to:

“Use the following command:

```
configure forwarding fabric hash [default |source-port | packet {algorithm [crc |
xor] | dynamic-mode [spray | eligibility | none]}] {slot slot-number}”
```

L2VPN Sharing Commands

ExtremeXOS Command Reference Guide, L2VPN commands

xos0062252

Remove the following commands:

- `configure l2vpn sharing hash-algorithm`
- `configure l2vpn sharing ipv4`

Platform availability for the following commands should be changed to the following:

- `enable l2vpn sharing`
- `disable l2vpn sharing`

Platform Availability:

This command is available only on the following switches

- Summit X460-G2
- Summit X670-G2
- Summit X670
- Summit X770
- BlackDiamond X8

Link Layer Discovery Protocol (LLDP)

In the *ExtremeXOS User Guide*, under **LLDP chapter > Configuring and Managing LLDP > Enable and Disable LLDP**

xos0061698

The existing text:

“LLDP is disabled on all ports by default.”

Should change to:

“LLDP is enabled on all ports by default.”

In the *ExtremeXOS Command Reference Guide* for the following commands:

- `enable lldp ports`
- `disable lldp ports`

The existing text:

“Default: Disabled”

Should change to:

“Default: Enabled”

Match Conditions Supported for OSPF Import Policy

In the *ExtremeXOS User Guide*, in OSPF and OSPFv3 chapters.

xos0063675

The following information should appear:

For OSPFv2 and OSPFv3, only “Network Layer Reachability Information” (NLRI) and “route origin” can be used as matching criteria in policy rules; using “next_hop” as a matching criteria is not supported. Any other policy attribute is not recognized and is ignored.

MIB

In the *ExtremeXOS User Guide*, in **Supported Standards, Protocols, and MIBs chapter > RFC 2613 (SMON) > monVlanIdStatsTable**.

xos0062086

The following variables and their comments should be removed under “Supported Variables” since they are not supported in ExtremeXOS.

- smonVlanIdStatsNUcastPkts
- smonVlanIdStatsNUcastOverflowPkts
- smonVlanIdStatsNUcastHCPkts
- smonVlanIdStatsNUcastOctets
- smonVlanIdStatsNUcastOverflowOctets
- smonVlanIdStatsNUcastHCOctets

MLAG PIM-SM *,G Forwarding Limitation

In the *ExtremeXOS User Guide*, under **MLAG > Multicast Over MLAG Configuration**.

xos0062552

Add the following note:



Note

It is recommended that for PIM-SM deployments, route to source must exist and that receivers should get the traffic from the SPT tree in the MLAG configurations. *,G forwarding in MLAG is not a recommended configuration.

NetLogin Limitation

In the *ExtremeXOS User Guide* under chapter **ExtremeXOS User Guide > Configuring Network Login > Configuring Network Login**

xos0062943

Add the following additional limitation:

“When using Netlogin MAC-based VLAN mode, moving a port as untagged from the pre-authentication VLAN to the post-authentication VLAN is not supported when both VLANs are configured with Protocol Filter IP.”

NetLogin Local Authentication

In the *ExtremeXOS User Guide*, under chapter **Network Login > Authentication Failure and Services Unavailable Handling > Dependency on authentication database order**

xos0063091

Remove the following section:

"For local authentication, if the user is not created in the local database, it is considered as service unavailable. If the user is configured but the password does not match, it is considered as an authentication failure."

In the *ExtremeXOS Command Reference Guide*, under the command `configure netlogin authentication service-unavailable vlan`

Remove the following section:

"For local authentication if the user entry is not present in the local database"

PoE Power Delivery

In the *ExtremeXOS User Guide*, under **POE > Power Delivery**.

xos0062239

Add the following note:

Note



In Summit X440 (PoE-capable) series switches, do not increase the power budget abruptly by a large extent.

When the power budget is suddenly, significantly increased, it exceeds the capabilities of the power supply unit, and the inline power state of the ports are disabled. If this occurs, reboot the switch to recover.

PVLAN + EAPS and PVLAN + STP Recommendations

In the *ExtremeXOS User Guide*, in **VLANs > Private VLANs > PVLAN Overview > PVLAN Limitations**

xos0063369

The following text:

"EAPS can only be configured on network VLAN ports (and not on subscriber VLAN ports). To support EAPS on the network VLAN, you must add all of the VLANs in the PVLAN to the EAPS ring.

STP can only be configured on network VLAN ports (and not on subscriber VLAN ports). To support STP on the network VLAN, you must add all of the VLANs in the PVLAN to STP."

Needs to be changed to:

"For PVLAN with STP implementation, irrespective of port translation configuration in Network VLAN, it is recommended to add both Network VLAN and all subscriber VLANs to STP.

For PVLAN with EAPS implementation, irrespective of port translation configuration in Network VLAN, it is recommended to add both Network VLAN and all subscriber VLANs to the EAPS ring."

VRRP Guidelines

In the *ExtremeXOS User Guide*, under **VRRP chapter > VRRP Guidelines**.

xos0056279, xos0062263

The VRRP guidelines should change to the following:

“The following guidelines apply to using VRRP:

- The maximum number of supported VRIDs per interface is seven.
- An interconnect link between VRRP routers should not be used, except when VRRP routers have hosts directly attached.
- A maximum of 128 VRID instances are supported on the router. This number can be extended up to 256 based on the license and hardware; refer to the release notes for the maximum limit.
- Up to seven unique VRIDs can be configured on the router. VRIDs can be re-used, but not on the same interface. When VRRP is configured for both IPv4 and IPv6, two unique VRIDs are consumed, though the same VRID is specified during VRRP instance creation. In other words, VRIDs used for each data protocol are counted separately.
- VRRP and other L2 redundancy protocols can be simultaneously enabled on the same switch.
- We do not recommend simultaneously enabling VRRP and ESRP on the same switch.
- When VRRP and BOOTP/DHCP relay are both enabled on the switch, the relayed BOOTP agent IP address is the actual switch IP address, not the virtual IP address. VRRP and ESRP cannot be configured on the same VLAN or port. This configuration is not allowed.
- RFC 5798 describes a situation where a master VRRP router takes on a duplicate IP address due to interaction with the duplicate address detection (DAD) feature. To prevent such duplicate addresses, the DAD feature is disabled whenever a VRRP router is configured for IPv6 or IPv4.
- A VRRP router instance can be configured with multiple IP addresses on the same subnet or on different subnets, provided that all virtual IP addresses match the subnet address of a VLAN on the switch. For example, if a host switch has VLAN IP addresses in the 1.1.1.x and 2.2.2.x subnets, then that VRRP router instance can contain virtual IP addresses in both those subnets as well. If a VRRP router instance is assigned priority 255, then the host router must own all the IP addresses assigned to the VRRP router instance. That is, each virtual IP address must match an IP address configured for a VLAN on the router.
- When a VRRPv2 instance spans routers using ExtremeXOS version 12.6 and earlier and routers using ExtremeXOS version 12.7 and later, routers using ExtremeXOS version 12.6 and earlier log packet-size warning messages.
- VRRP scaling numbers differs based on the license and hardware used; please refer the release notes for individual scaling limits.
- The maximum number of VIPs supported for a single VRRP instance is 255.”