



# ExtremeXOS Release Notes

*Software Version ExtremeXOS 15.6.4-Patch1-5*

Copyright © 2016 Extreme Networks, Inc. All Rights Reserved.

## Legal Notices

Extreme Networks, Inc., on behalf of or through its wholly-owned subsidiary, Enterasys 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 about Extreme Networks trademarks, go to:  
[www.extremenetworks.com/company/legal/trademarks/](http://www.extremenetworks.com/company/legal/trademarks/)

## Support

For product support, including documentation, visit:  
[www.extremenetworks.com/support/](http://www.extremenetworks.com/support/)

## Contact

Extreme Networks, Inc.  
145 Rio Robles  
San Jose, CA 19534  
Tel: +1 408-579-2800  
Toll-free: +1 888-257-3000

# Table of Contents

---

<b>Overview.....</b>	<b>6</b>
New and Corrected Features in ExtremeXOS 15.6 .....	6
Dynamic Host Configuration Protocol (DHCP) v6 RFC3315 Client .....	7
Supported Platforms .....	7
Limits .....	7
New CLI Commands .....	7
Changed CLI Commands .....	7
Python Scripting .....	8
Supported Platforms .....	8
Limitations .....	8
Changed CLI Commands .....	8
Inter-Virtual Router (VR) Routing for IPv4 Unicast Static Routes .....	8
Deployment Recommendations .....	9
Supported Platforms .....	9
Limitations .....	9
Changed CLI Commands .....	9
Link Aggregation (LAG) Configuration Change Does Not Disrupt Spanning Tree Protocol (STP) Operation .....	9
Supported Platforms .....	10
Limitations .....	10
Changed CLI Commands .....	10
Enable IProute Compression by Default .....	10
Supported Platforms .....	10
Changed CLI Commands .....	11
Layer 2 Protocol Tunneling (L2PT) and Filtering for Multi-protocol Label Switching (MPLS) .....	12
Supported Platforms .....	12
Limitations .....	12
Existing CLI Commands .....	13
Port-Specific VLAN Tagging .....	14
Supported Platforms .....	14
Limitations .....	15
Changed CLI Commands in ExtremeXOS 15.4.1 .....	15
ExtremeXOS Images for Summit X480 Series Switches .....	16
New Hardware Supported in ExtremeXOS 15.6 .....	17
Hardware Issues in ExtremeXOS 15.6 .....	17
Joint Interoperability Test Command (JITC) Compliance .....	18
ExtremeXOS Hardware and Software Compatibility Matrix .....	18
Upgrading to ExtremeXOS .....	19
Downloading Supported MIBs .....	19
Tested Third-Party Products .....	19
Tested RADIUS Servers .....	19
Tested Third-Party Clients .....	20
PoE Capable VoIP Phones .....	20

Extreme Switch Security Assessment .....	21
DoS Attack Assessment .....	21
ICMP Attack Assessment .....	21
Port Scan Assessment .....	21
Service Notifications .....	21
<b>Limits.....</b>	<b>22</b>
<b>Open Issues, Known Behaviors, and Resolved Issues .....</b>	<b>73</b>
Open Issues .....	73
Known Behaviors .....	76
Resolved Issues in ExtremeXOS 15.6.4-Patch1-5 .....	77
Resolved Issues in ExtremeXOS 15.6.4-Patch1-3 .....	78
Resolved Issues in ExtremeXOS 15.6.4 .....	81
Resolved Issues in ExtremeXOS 15.6.3-Patch1-9 .....	83
Resolved Issues in ExtremeXOS 15.6.3-Patch1-8 .....	85
Resolved Issues in ExtremeXOS 15.6.3-Patch1-5 .....	87
Resolved Issues in ExtremeXOS 15.6.3-Patch1-3 .....	89
Resolved Issues in ExtremeXOS 15.6.3 .....	93
Resolved Issues in ExtremeXOS 15.6.2-Patch1-6 .....	94
Resolved Issues in ExtremeXOS 15.6.2-Patch1-1 .....	97
Resolved Issues in ExtremeXOS 15.6.2 .....	99
Resolved Issues in ExtremeXOS 15.6.1 .....	104
<b>ExtremeXOS Documentation Corrections.....</b>	<b>109</b>
ACLs .....	110
ACL Egress Counters Limitation .....	110
ACL Policy Redirect .....	111
ACL Ports Limits .....	111
BOOTPrelay Not Supported in Virtual Routing and Forwarding (VRF) .....	112
ExtremeXOS User Guide .....	112
ExtremeXOS Command Reference .....	112
Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines ..	113
Configure IP-MTU VLAN Command Syntax Description .....	114
Command Reference .....	114
User Guide .....	114
Configure Sys-Recovery-level Slot Command Platform Availability .....	115
Debounce Commands .....	116
Configure stack-ports debounce time .....	116
Description .....	116
Syntax Description .....	116
Default .....	116
Usage Guidelines .....	116
Example .....	116
History: .....	116
Platforms Availability .....	116

---

Show stack-ports debounce .....	117
Description .....	117
Syntax Description .....	117
Default .....	117
Usage Guidelines .....	117
Example .....	117
History .....	117
Platform Availability .....	117
ELRP and QoS .....	118
ELRP on VPLS .....	119
L2VPN Sharing Commands .....	120
Link Aggregation (LAG) Limit for Multiprotocol Label Switching (MPLS) Terminated Packets .....	121
Link Layer Discovery Protocol (LLDP) .....	121
PoE Power Delivery .....	122
MLAG .....	122
MLAG PIM-SM *,G Forwarding Limitation .....	123
NetLogin Limitation .....	123
NetLogin Local Authentication .....	124
PVLAN + EAPS and PVLAN + STP Recommendations .....	125
Rate Limiting/Meters .....	126
Remote Mirroring .....	127
Routing Policies .....	128
Synchronize Command .....	129
TACACS Server .....	129
Unconfigure Switch Erase Command .....	131
VRRP Guidelines .....	132

# 1 Overview

---

These release notes document ExtremeXOS® 15.6.4-Patch1-5 which resolves software deficiencies.

This chapter contains the following sections:

- [New and Corrected Features in ExtremeXOS 15.6 on page 6](#)
- [ExtremeXOS Images for Summit X480 Series Switches on page 16](#)
- [Joint Interoperability Test Command \(JITC\) Compliance on page 18](#)
- [New Hardware Supported in ExtremeXOS 15.6 on page 17](#)
- [Hardware Issues in ExtremeXOS 15.6 on page 17](#)
- [ExtremeXOS Hardware and Software Compatibility Matrix on page 18](#)
- [Upgrading to ExtremeXOS on page 19](#)
- [Downloading Supported MIBs on page 19](#)
- [Tested Third-Party Products on page 19](#)
- [Extreme Switch Security Assessment on page 21](#)
- [Service Notifications on page 21](#)

## **New and Corrected Features in ExtremeXOS 15.6**

---

This section lists the new and corrected features supported in the ExtremeXOS 15.6 software:

- [Dynamic Host Configuration Protocol \(DHCP\) v6 RFC3315 Client on page 7](#)
- [Python Scripting on page 8](#)
- [Inter-Virtual Router \(VR\) Routing for IPv4 Unicast Static Routes on page 8](#)
- [Link Aggregation \(LAG\) Configuration Change Does Not Disrupt Spanning Tree Protocol \(STP\) Operation on page 9](#)
- [Enable IProute Compression by Default on page 10](#)
- [Layer 2 Protocol Tunneling \(L2PT\) and Filtering for Multi-protocol Label Switching \(MPLS\) on page 12](#)
- [Port-Specific VLAN Tagging on page 14](#)

## Dynamic Host Configuration Protocol (DHCP) v6 RFC3315 Client

The DHCPv6 for IPv6 enables DHCPv6 servers to pass configuration parameters such as IPv6 network addresses to a DHCPv6 client. The basic DHCPv6 client-server concept is similar to DHCP for IPv4. The DHCPv6 server and client exchange IPv6 messages using User Datagram Protocol (UDP.)

### Supported Platforms

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

### Limits

- Stateless Auto-configuration is not supported.
- Stateless DHCPv6 is not supported.
- Rapid-commit option is not supported.
- Temporary Address (IA\_TA) is not supported.
- Prefix Delegation (IA\_PD) is not supported.

### New CLI Commands

```
configure dhcp ipv6 client identifier-type [link-layer
{plus-time} | vendor-specific]
```

### Changed CLI Commands

Changes are bolded:

```
show dhcp-client {ipv4 | ipv6} state {<vlan>}
enable [bootp {ipv4}|dhcp {ipv4 | ipv6}] vlan [<vlan> |
all]
disable [bootp {ipv4}|dhcp {ipv4 | ipv6}] vlan [<vlan> |
all]
```

## Python Scripting

The `load script <script>` command can now be a Python v2.7.3 compatible script and may interact with the CLI and file system.

### Supported Platforms

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

### Limitations

Only allows a script to interact directly with the CLI interface for managing ExtremeXOS functionality.

### Changed CLI Commands

The `load script <script>` has been modified to accept Python v2.7.3 compatible scripts.

## Inter-Virtual Router (VR) Routing for IPv4 Unicast Static Routes

This feature supports IPv4 Inter-VR routing via static routes.

Virtual routers (VR) allow for separate Layer 3 routing domains. Each VR has its own routing table. VR in this context denotes any routable VR, and includes default VR, user VRs, VPN-VRFs, and non-VPN VRFs.

In some scenarios there is a need for communication between the isolated Layer 3 routing domains in a controlled way. For example, a university might decide to separate Layer 3 traffic on a per-department basis, but each department needs access to certain services provided by the payroll department. To accomplish this, you could install the routes related to those payroll services in the other departments' route tables, with the next hop interface corresponding to the payroll department's VR.



## Deployment Recommendations

To effectively deploy inter-VR routing using static routes, you need unique IP addresses for intended packet flows. Also note that during transient network conditions involving dynamic routing protocols, IP packets on outbound and/or return paths could match a less-specific static inter-VR route and be forwarded into another VR. Therefore, more-specific static inter-VR routes are recommended.

## Supported Platforms

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

## Limitations

This feature is not supported on MPLS, tunnels, or for multicast packets.

## Changed CLI Commands

The following command's syntax has not been changed, but will now accept that the egress VLAN name specified can belong to a different VR than the VR to which the route is added:

```
configure iproute add <ipNetmask> <gateway> {vlan  
<egress_vlan>} {vr <vr_name>}
```

## Link Aggregation (LAG) Configuration Change Does Not Disrupt Spanning Tree Protocol (STP) Operation

Currently, adding or deleting ports from a link aggregation group (LAG) produces the following error:

```
Warning: Any config on the master port is lost (STP, IGMP  
Filter, IGMP Static Group, MAC-Security, CFM, TRILL, etc.)
```

The current LAG implementation modifies the switch hardware VLAN interface tables by removing the entire LAG configuration and replaces it with the new LAG configuration. This destruction and recreation of virtual port interfaces may exhibit the characteristics of interface flapping and may cause traffic loss. This behavior is independent of any protocols configured on a LAG interface.

## Supported Platforms

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

## Limitations

- When aggregator membership changes on a LAG, both redistribution of flows, as well as some traffic loss, can occur. The implementation of this feature will not change this behavior.
- This enhancement does not affect the loss of configuration that occurs when executing the `enable sharing` and `disable sharing` commands.

## Changed CLI Commands

After executing either of the following commands there is no loss of configuration nor traffic:

```
configure sharing <port> add ports <port_list>
configure sharing <port> delete ports <port_list>
```

## Enable IProute Compression by Default

Compressed routes reduce the number of routes that are installed in the hardware routing tables. This improves packet forwarding performance when the switch uses hardware routing tables. The hardware routing tables have less storage space than the software, so compressed routes conserve resources and improve scaling.

This feature makes iproute compression on the default virtual router (VR), the default setting for new configurations. Additionally, this feature enables compression for user VR that are created on new and existing configurations. Both the IPv4 and IPv6 address family routing tables support this feature.

## Supported Platforms

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

## Changed CLI Commands

The `show configuration` command now shows compression as the default setting (in bold):

```
# show configuration "rtmgr" detail
#
# Module rtmgr configuration.
#
disable iproute sharing vr VR-Default
disable iproute ipv6 sharing vr VR-Default
disable iproute sharing vr VR-Mgmt
disable iproute ipv6 sharing vr VR-Mgmt
configure iproute priority mpls 20 vr VR-Default
configure iproute priority blackhole 50 vr VR-Default
configure iproute priority static 1100 vr VR-Default
configure iproute priority icmp 1200 vr VR-Default
. . .
enable iproute ipv4 compression
enable iproute ipv6 compression
```

The `show ipconfig (IP4)` command now shows compression as the default setting (in bold):

```
#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
...

```

The `show ipconfig ipv6` command now shows compression as the default setting (in bold):

```
#show ipconfig ipv6
      Route Sharing: Disabled
      Route Compression: Enabled
      ICMP Redirect for Fast Path: Disabled
      Max Shared Gateways: Current: 4 Configured: 4
...

```

## Layer 2 Protocol Tunneling (L2PT) and Filtering for Multi-protocol Label Switching (MPLS)

This feature was first introduced in ExtremeXOS 15.5.1. This revision of the feature adds MPLS support. Changes to the feature to provide MPLS support are marked with bold.

L2PT is useful for connecting remote switches across a service provider network. This feature allows control PDUs to be tunneled through the network, and provides a single STP domain for subscribers across the service provider network. Using tunneling the service provider network can be made transparent to the customer network.

Layer 2 PDU filtering allows a service provider to specify which Layer 2 PDUs are to be dropped at the ingress interface on a provider edge switch. Specified Layer 2 PDU frames are not transmitted over the service provider network.

### Supported Platforms

- BlackDiamond X8 and BlackDiamond 8800<sup>1</sup> series switches
- Summit X770, X670, X670-G2, X480<sup>1</sup>, X460, X460-G2, and X440 series switches
- E4G-200 and E4G-400 cell site routers

### Limitations

Existing limits and new ones for ExtremeXOS 15.6 (in bold):

- L2PT and protocol filtering is implemented in software, thus limiting the number of frames that may be filtered or tunneled.
- Both L2PT and protocol filtering may be configured only through command line interface. Configuration through SNMP/XML is not supported.
- **If L2PT configurations are made on pseudowires, these configurations are lost after restarting the MPLS process unless the L2PT process is also restarted.**
- **If L2PT configurations are made on a VPLS or VPWS service, dot1p tag inclusion must be enabled on the VPLS/VPWS.**
- When tunneling protocols which are point-to-point in nature, you must ensure that there are only two tunnel endpoints for the protocol.

---

1. L2PT over VPLS/VPWS is not supported on Summit X480 series switches and in BlackDiamond 8800 XL modules.

- If a protocol that is configured to be tunneled on a service interface cannot be uniquely identified by its DA and EtherType, then all packets with the same DA and EtherType of the protocol being tunneled but are not really PDUs of the protocol are slow path forwarded.
- Tagged protocol PDUs cannot be tunneled over VLANs. Tagged protocol PDUs can be tunneled only over VMANs (the VMAN can be the service VMAN for a VPLS/VPWS service or a standalone VMAN). Untagged protocol PDUs can be tunneled over both VLANs and VMANs (the VLAN/VMAN can be standalone or be the service VMAN for a VPLS/VPWS service).
- Untagged protocol PDUs cannot be bypassed if the ingress port is an untagged VMAN port with a default CVID. Untagged protocol PDUs may be bypassed if the ingress port is an untagged VMAN port without a default CVID.
- L2PT is not supported on VLAN ports that have a port specific tag.
- **In VPLS, only full-mesh configuration is supported for L2PT.**

### Existing CLI Commands

- `configure protocol filter <filter_name> [add | delete] dest-mac <mac_address> {[etype | llc | snap] <hex>} {field offset <offset> value <value> {mask <mask>}}`
- `configure l2pt encapsulation dest-mac <mac_address>`
- `show l2pt`
- `create l2pt profile <profile_name>`
- `delete l2pt profile <profile_name>`
- `configure l2pt profile <profile_name> add protocol filter <filter_name> {action [tunnel {cos <cos>} | encapsulate | none]}`
- `configure l2pt profile <profile_name> delete protocol filter <filter_name>`
- `show l2pt profile {<profile_name>}`
- `configure [vlan | vman] <vlan_name> ports <port_list> l2pt profile [none | <profile_name>]`
- `show [vlan | vman] <vlan_name> {ports <port_list>} l2pt {detail}`
- `clear l2pt counters {[vlan | vman] <vlan_name> {ports <port_list>}}`
- `configure ports [<port_list> | all] protocol filter [none | <filter_name>]`

- `show ports [<port_list> | all] protocol filter {detail}`
- `clear counters ports {<port_list> | all} protocol filter`
- `create protocol {filter} <filter_name>`
- `delete protocol {filter} <filter_name>`
- `configure protocol {filter} <filter_name> [add | delete] [etype | llc | snap] <hex> {[etype | llc | snap] <hex>} {[etype | llc | snap] <hex>} {[etype | llc | snap] <hex>}`
- `show protocol {filter} {<filter_name>} {detail}`
- `configure {vlan} <vlan_name> protocol {filter} <filter_name>`

## Port-Specific VLAN Tagging

This feature allows bridging of frames with different VLAN IDs that have been trunked by third-party equipment. Frames with different VLAN IDs are treated as belonging to the same VLAN.

Different VLAN IDs can be on different ports or on the same ports. Different VLAN IDs are used to accept packets to the VLAN, and the right VLAN ID is used when forwarding the frames out. The same VLAN ID on different ports can be associated with different VLANs. Forwarding can also be done over L2VPN when the switches are connected with pseudowires.

This feature was first introduced in ExtremeXOS 15.4.1. This revision of the feature allows VLANs which have port-specific tags to be an EAPS-protected VLAN, provided that the port-specific tag “attachments” is not configured on the EAPS (primary and secondary) ring ports. Changes to the feature to support this enhancement are marked with bold.

### *Supported Platforms*

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

## Limitations

- Protocols are not supported on port-specific VLAN tags. They are prevented from being configured. **VLANs which have port-specific tags may be EAPS-protected VLANs, provided that the port-specific tag “attachments” is not configured on the EAPS (primary and secondary) ring ports.**
- VMANs are not supported.
- Configuration input and output are only through the CLI.
- IP unicast/multicast forwarding is not supported.
- Within a single broadcast domain (VLAN) there is only one MAC address. When a MAC address is learned on different tag (on the same port or on a different port), it is considered a MAC move.
- Multicast is not supported. Internet Group Management Protocol (IGMP) snooping must be disabled.
- VLAN translation functionality is essentially provided by VLAN bridging tag, so VLAN bridging tag cannot be part of VLAN translation (either as a translation VLAN, or a member VLAN).
- VLANs with a port-specific tag cannot be part of a private VLAN.
- Remote-mirroring is not supported on VLANs with port-specific tags.

### *Changed CLI Commands in ExtremeXOS 15.4.1*

- `configure {vlan} vlan_name add ports [port_list | all] {tagged {<tag>} | untagged} {{stpd} stpd_name} {dot1d | emistp | pvst-plus}}`
- `configure {vlan} vlan_name delete ports [all | port_list {tagged <tag>}]`
- `create fdbentry <mac_addr> vlan <vlan_name> [ports <port_list> {tagged <tag>} | blackhole]`
- `configure ports port_list {tagged <tag>} vlan vlan_name [limit-learning number {action[blackhole | stop-learning]} | lock-learning | unlimited-learning | unlocklearning]`

Additionally, the output of the following commands is modified to show VLAN bridging information:

- `show vlan`
- `show port info detail`
- `show fdb`

## ExtremeXOS Images for Summit X480 Series Switches

Due to additional functionality and new platforms supported, the ExtremeXOS 15.6 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 1: Summit X480 Series Switches Software Image Files**

	Main Install image	Diagnostic image
<b>Content</b>	All Summit X480 content (except diagnostics)	Summit X480 diagnostics
<b>File Name</b>	summitX480-15.6.xx.yy.xos	summitX480-15.6.xx.yy-diagnostics.xmod
<b>File Type</b>	Standard ExtremeXOS image	XMOD image
<b>Installation Notes</b>	<ul style="list-style-type: none"> <li>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.</li> <li>The Summit XMODs, such as SSH can be used with the summitX480 ExtremeXOS image.</li> </ul>	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-15.6.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-15.6.xx.yy.xos` file image, and this image installs the diagnostics capability.

---

## New Hardware Supported in ExtremeXOS 15.6

---

This section lists the new hardware supported in ExtremeXOS 15.6:

- Summit X670-G2 series switches:

X670-G2-48X-4q and X670-G2-72X

- Summit X460-G2 series switches:

X460-G2-24t-10GE4, X460-G2-48t-10GE4, X460-G2-24p-10GE4, X460-G2-48p-10GE4, X460-G2-24x-10GE4, X460-G2-48x-10GE4, X460-G2-24t-GE4, X460-G2-48t-GE4, X460-G2-24p-GE4, and X460-G2-48p-GE4

**NOTE**

On the Summit X670-G2 series switches, the Management Port LED does not work at 10M link speed (LED is off in 10M mode). The 1G and 100M LEDs do function properly.

---

**NOTE**

The Summit X460-G2 series switches do not support half duplex.

---

## Hardware Issues in ExtremeXOS 15.6

---

The E4G-200 cell site router front panel alarm DB15 connector capabilities are not currently supported.

---

## 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.

---

## ExtremeXOS Hardware and Software Compatibility Matrix

---

The *ExtremeXOS Hardware and Software Compatibility Matrix* provides information about the minimum version of ExtremeXOS software required to support BlackDiamond and Summit switches, as well as SFPs, XENPAKs, XFPs, and other pluggable interfaces.

The latest version of the *ExtremeXOS Hardware and Software Compatibility Matrix* can be found at:

[www.extremenetworks.com/documentation](http://www.extremenetworks.com/documentation)

---

## Upgrading to ExtremeXOS

---

For instructions about upgrading ExtremeXOS software, see the “Software Upgrade and Boot Options” chapter in the Basic Switch Operation volume of 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." is displayed. 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

---

## Downloading Supported MIBs

---

The Extreme Networks MIBs are located on the eSupport website under **Download Software Updates**, located at:

[https://esupport.extremenetworks.com/eservice\\_enu/start.swe?SWECmd=Start&SWEHo=esupport.extremenetworks.com](https://esupport.extremenetworks.com/eservice_enu/start.swe?SWECmd=Start&SWEHo=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 ExtremeXOS 15.5.

### 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:

<http://www.extremenetworks.com/support/service-notification-form>

# 2 Limits

---

This chapter summarizes the supported limits in ExtremeXOS 15.6.4-Patch1-5.

**Table 2** 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 2** 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 2** 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 2: Supported Limits

Metric	Product	Limit
<b>AAA (local)</b> —maximum number of admin and local user accounts.	All platforms	8
<b>Access lists (meters)</b> —maximum number of meters.	BlackDiamond 8000 series e-series, group of 24 ports c-series  BlackDiamond 8900 series 8900-10G24X-c, group of 12 ports  8900 xl-series, 8900-G96T-c  8900-40G6X-xm  BlackDiamond X8 a-series modules  BlackDiamond X8-100G4X modules  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 X770, X670-G2, X460-G2	512 2,048 ingress, 256 egress  1,024 ingress, 256 egress  4,096 ingress, 512 egress  512 ingress 512 egress  512 ingress, 512 egress  8,192 ingress, 1,024 egress  1,024 ingress 256 egress  512 ingress  2,048 ingress, 256 egress  4,096 ingress, 512 egress  512 ingress 512 egress  1,024 ingress, 512 egress
<b>Access lists (policies)</b> —suggested maximum number of lines in a single policy file.	All platforms	300,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Access lists (policies)</b> —maximum number of rules in a single policy file. <sup>a</sup>	BlackDiamond 8000 series c-series, group of 24 ports	4,096 ingress, 512 egress
	e-series, group of 24 ports	1,024 ingress
	BlackDiamond 8900 8900-10G24X-c modules, group of 12 ports	2,048 ingress, 512 egress
	8900-G96T-c modules, group of 48 ports	8,192 ingress, 1,024 egress
	8900 xl-series 8900-40G6X-xm	61,440 (up to) 2,048 ingress, 1,024 egress
	BlackDiamond X8 a-series modules	2,048 ingress, 1,024 egress
	BlackDiamond X8-100G4X modules	8,192 ingress, 1,024 egress
	Summit X440, X430 group of 24 ports	1,024 ingress
	Summit X460, E4G-400	4,096 ingress, 512 egress
	Summit X480	(up to) 61,440 ingress, 1,024 egress
	Summit X670 VIM4-40G4x	2,048 ingress 1,024 egress
	Summit X480	8,192 ingress/ 1,024 egress
	Summit X480 VIM3-40G4X	2048 ingress 1024 egress
	Summit X770, X670-G2, X460-G2 E4G-200	4,096 ingress 1,024 egress 2,048 ingress/ 512 egress



Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Access lists (slices)</b> —number of ACL slices.	BlackDiamond 8000 series c-series, group of 48 ports	16
	BlackDiamond 8900 series 8900-10G24X-c modules, group of 12 ports	12 ingress, 4 egress
	8900-G96T-c modules, group of 48 ports	16 ingress, 4 egress
	8900 xl-series	17 <sup>b</sup>
	8900-40G6X-xm	10 ingress, 4 egress
	BlackDiamond X8 a-series modules	10 ingress, 4 egress
	BlackDiamond X8-100G4X modules	16 ingress, 4 egress
	E4G-200	8 ingress, 4 egress
	Summit X440, X430	4 ingress
	Summit X460, E4G-400, X460-G2	16 ingress, 4 egress
	Summit X480	17 <sup>b</sup> ingress, 4 egress
	Summit X670 VIM4-40G4x	10 ingress, 4 egress
	Summit X480 VIM3-40G4X	10 ingress, 4 egress
Summit X770, X670-G2	12 ingress 4 egress	
<b>AVB (audio video bridging)</b> —maximum number of active streams. <b>NOTE:</b> * It is recommended that you do not use on more than 8 ports on this switch.	Summit X440, X460, X460-G2	1,024
	Summit X670, X670, X670-G2	4,096
	Summit X430	100*
<b>BFD sessions</b> —maximum number of BFD sessions.	All platforms (default timers—1 sec)	512
	BlackDiamond X8 and 8800 (minimal timers—50 msec)	10 <sup>c</sup>
	All Summits (minimal timers—100 msec)	10 <sup>c</sup>

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>BGP (aggregates)</b> —maximum number of BGP aggregates.	All platforms (except E4G-200, X430, and X440) with Core license or higher	256
<b>BGP (networks)</b> —maximum number of BGP networks.	All platforms (except X430, and X440) with Core license or higher	1,024
	BlackDiamond X8 series	1,024
<b>BGP (peers)</b> —maximum number of BGP peers. <b>NOTE:</b> * With default keepalive and hold timers.	BlackDiamond X8 series	512
	BlackDiamond 8000 series	512
	BlackDiamond xl-series	512
	All Summits, except X480, X440, X430	128*
	E4G-400	128*
Summit X480	512	
<b>BGP (peer groups)</b> —maximum number of BGP peer groups.	BlackDiamond 8900 series	128
	BlackDiamond 8800	64
	BlackDiamond X8 series	128
	Summit X480	128
	Summit X770, X670-G2, X670v-48t, X670, X460-G2, X460 (with Core license or higher)	64
<b>BGP (policy entries)</b> —maximum number of BGP policy entries per route policy.	All platforms (except E4G-200, X430, and X440) with Core license or higher	256
<b>BGP (policy statements)</b> —maximum number of BGP policy statements per route policy.	All platforms (except E4G-200, X430, and X440) with Core license or higher	1,024
<b>BGP multi-cast address-family routes</b> —maximum number of multi-cast address-family routes.	BlackDiamond 8900 xl-series	524,256 (up to) <sup>b</sup>
	BlackDiamond X8 series	25,000
	Summit X460, X460-G2, X670, X670-G2, X770,	25,000
	Summit X480	524,256 (up to) <sup>b</sup>
	E4G-400	25,000
<b>BGP (unicast address-family routes)</b> —maximum number of unicast address-family routes.	BlackDiamond 8900 xl-series	524,256 (up to) <sup>b</sup>
	BlackDiamond X8 series	25,000
	Summit X460, X460-G2, X670, X670-G2, X770	25,000
	Summit X480	524,256 (up to) <sup>b</sup>
	E4G-400	25,000

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>CFM</b> —maximum number of CFM up end points. <b>NOTE:</b> With Advanced Edge license or higher.	BlackDiamond 8000 series	32
	BlackDiamond X8 series	32
	Summit series	32
<b>CFM</b> —maximum number of CFM down end points. <b>NOTE:</b> With Advanced Edge license or higher.	BlackDiamond 8000 series	32
	BlackDiamond X8 series	32
	Summit series X460, E4G-200, E4G-400 (non-load shared ports)	256 (non-load shared ports) 32 (load shared ports)
	All other platforms	32
<b>CFM</b> —maximum number of CFM remote end points per up/down end point. <b>NOTE:</b> With Advanced Edge license or higher.	All platforms	2,000
<b>CFM</b> —maximum number of dot1ag ports. <b>NOTE:</b> With Advanced Edge license or higher.	All Summits, except X430	128
<b>CFM</b> —maximum number of CFM segments. <b>NOTE:</b> With Advanced Edge license or higher.	All platforms	1,000
<b>CFM</b> —maximum number of MIPs. <b>NOTE:</b> With Advanced Edge license or higher.	All platforms	256
<b>CLEAR-Flow</b> —total number of rules supported. The ACL rules plus CLEAR-Flow rules must be less than the total number of supported ACLs.	Summit X440, X430	1,024
	Summit X670	2,048
	Summit X460, X460-G2, X770, X670-G2	4,096
	Summit X480	8,192
	E4G-200	2,048
	E4G-400	4,094
<b>Data Center Bridging eXchange (DCBX) protocol Type Length Value (TLVs)</b> —maximum number of DCBX application TLVs.	All platforms	8
<b>Dynamic ACLs</b> —maximum number of ACLs processed per second. <b>NOTE:</b> Limits are load dependent.	Summit X480, X670 with 50 DACLs with 500 DACLs	10 5

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>EAPS domains</b> —maximum number of EAPS domains. <b>NOTE:</b> An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains.	BlackDiamond 8000 series	64
	BlackDiamond X8 series	64
	Summit series (except X430), E4G-200, E4G-400	32
	Summit X430	4
<b>EAPsv1 protected VLANs</b> —maximum number of protected VLANs.	BlackDiamond 8000 series	2,000
	BlackDiamond X8 series	2,000
	Summit series, E4G-200, E4G-400	1,000
<b>EAPsv2 protected VLANs</b> —maximum number of protected VLANs.	BlackDiamond 8000 series	2,000
	BlackDiamond X8 series	2,000
	All Summits (except X430, X440), E4G-200, E4G-400	500
		500
<b>ELSM (vlan-ports)</b> —maximum number of VLAN ports.	BlackDiamond 8000 series	5,000
	BlackDiamond X8 series	5,000
	All Summits, E4G-200, E4G-400	5,000
<b>ERPS domains</b> —maximum number of ERPS domains without CFM configured.	BlackDiamond 8806 series	32
	BlackDiamond X8 series	32
	Summit series (except X430), E4G-200, E4G-400	32
	Summit X430	4
<b>ERPS domains</b> —maximum number of ERPS domains with CFM configured.	BlackDiamond 8806 series	16
	BlackDiamond X8 series	16
	Summit X440, X770, X670, X670-G2, X480, X460-G2	16
	Summit X460	32
	Summit X430	4
	E4G-200, E4G-400	32
<b>ERPSv1 protected VLANs</b> —maximum number of protected VLANs.	BlackDiamond 8806 series	2,000
	BlackDiamond X8 series	2,000
	All Summits, E4G-200, E4G-400	1,000
<b>ERPSv2 protected VLANs</b> —maximum number of protected VLANs.	BlackDiamond 8806 series	2,000
	BlackDiamond X8 series	2,000
	All Summits (except X430), E4G-200, E4G-400	500
<b>ESRP groups</b> —maximum number of ESRP groups.	All platforms	7
<b>ESRP domains</b> —maximum number of ESRP domains.	All platforms	64

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>ESRP VLANs</b> —maximum number of ESRP VLANs.	BlackDiamond 8000/8900 series	1,000
	BlackDiamond X8	2,048
	All Summits	1,000
	E4G-200, E4G-400	1,000
<b>ESRP (maximum ping tracks)</b> —maximum number of ping tracks per VLAN.	All platforms (except Summit X430)	8
<b>ESRP (IP route tracks)</b> —maximum IP route tracks per VLAN.	All platforms (except Summit X430)	8
<b>ESRP (VLAN tracks)</b> —maximum number of VLAN tracks per VLAN.	All platforms (except Summit X430)	1
<b>Forwarding rate</b> —maximum L2/L3 software forwarding rate.	BlackDiamond 8000 series	10,000 pps
	BlackDiamond X8 series	20,000 pps
	Summit X770	13,842 pps
	Summit X670-G2	29,028 pps
	Summit X670	14,829 pps
	Summit X480	14,509 pps
	Summit X460-G2	29,037 pps
	Summit X460	5,222 pps
	Summit X440	5,418 pps
E4G-200	8,718 pps	
E4G-400	5,536 pps	
<b>FDB (blackhole entries)</b> —maximum number of unicast blackhole FDB entries.	BlackDiamond 8800 c-series	32,000
	BlackDiamond 8000 e-series	8,000
	BlackDiamond 8900 series	
	8900 c-series	32,000
	8900 xl-series	524,288 (up to) <sup>b</sup>
	8900-40G6X-xm	128,000
	BlackDiamond X8 a-series modules	128,000
	BlackDiamond X8-100G4X modules	384,000
	E4G-200, E4G-400	32,000
	Summit X440, X430	16,000
	Summit X480	524,288 (up to) <sup>b</sup>
	Summit X460	32,000
	Summit X460-G2	49,152 <sup>d</sup>
Summit X670 VIM4-40G4x, X480 VIM3-40G4X	128,000	
Summit X770, X670-G2	294,912 <sup>d</sup>	
Summit X670, X670v-48t	130,000 <sup>d</sup>	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>FDB (blackhole entries)</b> — maximum number of multi- cast blackhole FDB entries.	BlackDiamond 8000 series	1,024
	BlackDiamond X8 series	1,024
	Summit X480, X460-G2, X460, X440, X430	1,024
	Summit X770, X670, X670-G2, X670v-48t, X480 VIM3-40G4X	4,096
	E4G-200, E4G-400	1,024
<b>FDB (maximum L2 entries)</b> — maximum number of MAC addresses.	BlackDiamond 8000 c-series	32,768 <sup>e</sup>
	BlackDiamond 8000 e-series	8,192 <sup>e</sup>
	BlackDiamond 8000 (system), except 8900 xl-series	128,000 <sup>e</sup>
	BlackDiamond 8900 xl-series	524,488 (up to) <sup>b</sup>
	BlackDiamond X8 a-series modules	128,000 <sup>e</sup>
	BlackDiamond X8-100G4X modules	384,000 <sup>e</sup>
	E4G-200, E4G-400	32,000 <sup>e</sup>
	Summit X440, X430	16,000 <sup>e</sup>
	Summit X480	524,488 (up to) <sup>b</sup>
	VIM3-40G4X module	128,000 <sup>e</sup>
	Summit X460	32,000 <sup>e</sup>
	Summit X460-G2	96,000 <sup>d e</sup>
	Summit X670	128,000 <sup>e</sup>
Summit X770	294,912 <sup>d e</sup>	
Summit X670-G2	228,000	
<b>FDB (Maximum L2 entries)</b> — maximum number of multi- cast FDB entries.	BlackDiamond X8	1,024
	BlackDiamond 8800	1,024
	Summit X770, X670, X670-G2	4,096
	Summit X480, X460, X460-G2, X430, X440	1,024
	E4G-200, E4G-400	1,024
<b>FIP Snooping VLANs</b>	BlackDiamond X8	768
	BlackDiamond 8800 (8900-40G6X- c only)	
	Summit X670	
<b>FIP Snooping Virtual Links (FPMA mode) per port group</b>	BlackDiamond X8	1,908
	BlackDiamond 8800 (8900-40G6X- c only)	
	Summit X670	

Table 2: Supported Limits (Continued)

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



Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Identity management</b> —recommended number of ACL entries per identity. <b>NOTE:</b> Number of ACLs per identity based on system ACL limitation.	All platforms.	20
<b>Identity management</b> —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
<b>IGMP sender</b> —maximum number of IGMP senders per switch (IP multi-cast compression disabled). <sup>l</sup> <b>NOTE:</b> Assumes source-group-vlan 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 X8-100G4X E4G-200, E4G-400 Summit X440 Summit X480 Summit X460 Summit X460-G2 Summit X670 VIM4-40G4x Summit X670-G2 Summit X770 Summit X430	2,048 <sup>f</sup> 500 <sup>g</sup> 2,048 <sup>f</sup> 4,096 <sup>f</sup> 3,000 <sup>g</sup> 4,096 <sup>f</sup> 4,096 <sup>h</sup> 16,384 <sup>d</sup> 2,048 64 4,096 2,048 4,096  3,000 <sup>g</sup> 4,096 <sup>g</sup> 4,096 64

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p><b>IGMP sender</b>—maximum number of IGMP senders per switch (IP multi-cast compression enabled).<sup>1</sup></p> <p><b>NOTE:</b> Assumes source-group-vlan mode.</p> <p>For additional limits, see:</p> <ul style="list-style-type: none"> <li>• <b>Layer-2 IPMC forwarding caches</b>—(IGMP/MLD/PIM snooping) in <b>mac-vlan mode</b>. on page 46</li> <li>• <b>Layer-2 IPMC forwarding caches</b>— (IGMP/MLD/PIM snooping) in <b>mixed-mode</b>. on page 47</li> </ul>	BlackDiamond 8800 c-series	2,048 <sup>g</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900-10G24X-c modules	2,048 <sup>g</sup>
	BlackDiamond 8900-G96T-c modules	4,096 <sup>g</sup>
	BlackDiamond 8900-40G6X-xm	3,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	4,096 <sup>g</sup>
	BlackDiamond X8 a-series modules	4,096 <sup>b</sup>
	BlackDiamond X8-100G4X modules	16,384 <sup>g</sup>
	E4G-200	3,000 <sup>g h</sup>
	E4G-400	6,000 <sup>g h</sup>
	Summit X440	192 <sup>g</sup>
	Summit X460	6,000 <sup>g</sup>
	Summit X460-G2	21,000 <sup>g</sup>
	Summit X480	12,000 <sup>g</sup>
	Summit X670	3,000 <sup>g</sup>
VIM4-40G4x	12,000 <sup>b</sup>	
Summit X770, X670-G2	66,500 <sup>g</sup>	
Summit X430	192	
<p><b>IGMP snooping per VLAN filters</b>—maximum number of VLANs supported in per-VLAN IGMP snooping mode.</p>	BlackDiamond 8800 c-series	2,000
	BlackDiamond 8000 e-series	448
	BlackDiamond 8900 c-series	1,000
	BlackDiamond 8900 xl-series	4,000
	BlackDiamond 8900-40G6X-xm	1,000
	BlackDiamond X8 a-series modules	1,000
	BlackDiamond X8-100G4X modules	4,000
	E4G-200, E4G-400	1,000
	Summit X440	448
	Summit X460, X670, X440	1,000
	Summit X460-G2	1,350
	Summit X480	4,000
	Summit X770	2,000
Summit X670-G2	1,500	
<p><b>IGMPv1/v2 SSM-map entries</b>—maximum number of IGMPv1/v2 SSM mapping entries.</p>	All platforms	500
<p><b>IGMPv1/v2 SSM-MAP entries</b>—maximum number of sources per group in IGMPv1/v2 SSM mapping entries.</p>	All platforms	50

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IGMPv2 subscriber</b> —maximum number of IGMPv2 subscribers per port. <sup>m</sup>	BlackDiamond 8800 c-series	2,000
	BlackDiamond 8900 c-series	2,000
	BlackDiamond X8 series	2,000
	Summit X430, X460, E4G-200, E4G-400, X440	1,000
	Summit X480, X670, X670v-48t	2,000
	Summit X770, X670-G2, X460-G2	4,000
<b>IGMPv2 subscriber</b> —maximum number of IGMPv2 subscribers per switch. <sup>m</sup>	BlackDiamond 8800 c-series	20,000
	BlackDiamond 8900 c-series	20,000
	BlackDiamond X8 series	20,000
	Summit X430, X440, E4G-200	10,000
	Summit X460, X460-G2, X480, X670, E4G-400, X670v-48t	20,000
	Summit X770, X670-G2	30,000
<b>IGMPv3 maximum source per group</b> —maximum number of source addresses per group.	All platforms	250
<b>IGMPv3 subscriber</b> —maximum number of IGMPv3 subscribers per port. <sup>m</sup>	BlackDiamond 8800 e-series	1,000
	BlackDiamond 8800 c-series	2,000
	BlackDiamond 8900 series	5,000
	BlackDiamond X8 series	3,000
	Summit X480, X670, X670v-48t, E4G-200, X440	1,000
	Summit X770, X670-G2, X460-G2 Summit X460, E4G-400	4,000 2,000
<b>IGMPv3 subscriber</b> —maximum number of IGMPv3 subscribers per switch. <sup>m</sup>	BlackDiamond 8800 e-series	10,000
	BlackDiamond 8800 c-series	20,000
	BlackDiamond 8900 series	30,000
	BlackDiamond X8 series	30,000
	Summit X670, X670v-48t, X480, E4G-200, X440	10,000
	Summit X460, E4G-400	20,000
	Summit X770, X670-G2 Summit X460-G2	30,000 24,000
<b>IP ARP entries in software</b> —maximum number of IP ARP entries in software. <b>NOTE:</b> May be limited by hardware capacity of FDB (maximum L2 entries).	Summit X770	131,072 (up to) <sup>i</sup>
	BlackDiamond X8-100G4X modules	229,374 (up to) <sup>i</sup>
	Summit X670-G2	131,072
	Summit X670, X480, X460, X440, X430	20,480
	Summit X460-G2	57,344 (up to) <sup>i</sup>
	E4G-200, E4G-400	20,480

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IP ARP entries in software with distributed mode on</b> —maximum number of IP ARP entries in software with distributed mode on.	BlackDiamond 8000 series with 8900-MSM128 or MSM-48c, and only 8900 xl-series I/O modules	260,000
	BlackDiamond 8000 series with any I/O modules that are not 8900 xl-series	100,000
	BlackDiamond X8 series	28,000
	All other platforms	N/A
<b>IPv4 ARP entries in hardware with distributed mode on</b> —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 <sup>b</sup>
	Per BlackDiamond 8900-G48X-xl or 8900-G48T-xl, up to 130,000 per system	16,250 <sup>b</sup>
	Per BlackDiamond 8000 c-series, up to 18,000 per system	8,000
	BlackDiamond 8900-40G6X-xm, up to 22,000 per system	8,000
	BlackDiamond X8 series, up to 28,000 per system	12,000
All other platforms	N/A	
<b>IPv4 ARP entries in hardware with minimum LPM routes</b> —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 <sup>g</sup>
	BlackDiamond 8900 xl-series	16,000
	BlackDiamond X8 a-series	16,000
	BlackDiamond X8-100G4X modules	182,000(up to) <sup>i</sup>
	E4G-200	8,000
	E4G-400	16,000
	Summit X440	412
	Summit X670, X480 VIM3-40G4X	8,000
	Summit X460, X480	16,000
	Summit X460-G2	50,000 (up to) <sup>i</sup>
	Summit X770, X670-G2	108,000(up to) <sup>i</sup>

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IPv4 ARP entries in hardware with maximum LPM routes</b> —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 <sup>g</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 xl-series	12,000 <sup>g</sup>
	BlackDiamond X8 a-series	12,000 <sup>g</sup>
	BlackDiamond X8-100G4X modules	172,000 (up to) <sup>i</sup>
	E4G-200	6,000 <sup>g</sup>
	E4G-400	12,000 <sup>g</sup>
	Summit X440	380
	Summit X460, X480	12,000 <sup>g</sup>
	Summit X670, X480 VIM3-40G4X	6,000 <sup>g</sup>
	Summit X770, X670-G2	98,000 (up to) <sup>i</sup>
<b>IP flow information export (IPFIX)</b> —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 X8-100G4X 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
<b>IPv4 remote hosts in hardware with zero LPM routes</b> —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 <sup>g</sup>
	BlackDiamond 8000 e-series	1,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	40,000 <sup>b</sup>
	BlackDiamond 8900-40G6X-xm	22,000 <sup>g</sup>
	BlackDiamond X8 a-series	28,000 <sup>g</sup>
	BlackDiamond X8-100G4X modules	311,000 (up to) <sup>i</sup>
	E4G-200	18,000 <sup>g</sup>
	E4G-400	20,000 <sup>g</sup>
	Summit X440	448
	Summit X460	20,000 <sup>g</sup>
	Summit X460-G2	73,000 <sup>i</sup>
Summit X480	40,000 <sup>b</sup>	
Summit X670, X480 VIM3-40G4X	22,000 <sup>g</sup>	
Summit X770, X670-G2	176,000 (up to) <sup>i</sup>	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IPv4 routes</b> —maximum number of IPv4 routes in software (combination of unicast and multi-cast routes).	BlackDiamond 8900 xl-series with 8900-MSM128 or MSM-48c	524,256 (up to) <sup>b</sup>
	All other BlackDiamond 8000 series hardware	25,000
	BlackDiamond X8 series	25,000
	Summit X440	256
	Summit X460, X670, X770, X670-G2, X460-G2	25,000
	Summit X480	524,256 (up to) <sup>b</sup>
<b>IPv4 routes (LPM entries in hardware)</b> — number of IPv4 routes in hardware.	E4G-200, E4G-400	25,000
	BlackDiamond 8800 c-series	12,000
	BlackDiamond 8000 e-series	480
	BlackDiamond 8900 xl-series	524,256 (up to) <sup>b j</sup>
	BlackDiamond 8900-40G6X-xm	16,000 <sup>e</sup>
	BlackDiamond X8 series	16,000 <sup>e</sup>
	E4G-200, E4G-400	12,000
	Summit X440	32
	Summit X460, X460-G2	12,000
	Summit X480	524,256 (up to) <sup>b j</sup>
<b>IPv6 addresses on an interface</b> — maximum number of IPv6 addresses on an interface.	Summit X480 VIM3-40G4X	16,000 <sup>i</sup>
	Summit X670	12,000
	Summit X770, X670-G2	16,000
	All platforms	255
<b>IPv6 addresses on a switch</b> — 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	2,048

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IPv6 host entries in hardware</b> —maximum number of IPv6 neighbor entries in hardware.	BlackDiamond 8800 c-, xm-series	3,000 <sup>g</sup>
	BlackDiamond 8000 e-series	250 <sup>g</sup>
	BlackDiamond 8900-10G24X-c modules	2,000 <sup>g</sup>
	BlackDiamond 8900-G96T-c modules	4,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	8,192 (up to) <sup>b g</sup>
	BlackDiamond X8 a-series	3,000 <sup>g</sup>
	BlackDiamond X8-100G4X modules	49,000 <sup>g</sup>
	E4G-200	2,000 <sup>g</sup>
	E4G-400	3,000 <sup>g</sup>
	Summit X440	192 <sup>g</sup>
	Summit X460, X670, X480 VIM3-40G4X	3,000 <sup>g</sup>
	Summit X770, X670-G2	36,750 <sup>g</sup>
	Summit X480, X670v-48t	6,000 <sup>g</sup>
Summit X460-G2	22,000 <sup>g</sup>	
<b>IPv6 routes (LPM entries in hardware)</b> —maximum number of IPv6 routes in hardware.	BlackDiamond 8800 c-series	6,000
	BlackDiamond 8000 e-series	240
	BlackDiamond 8900 xm-series	8,000
	BlackDiamond 8900 xl-series	245,760 (up to) <sup>b</sup>
	BlackDiamond X8 series	8,000
	E4G-200, E4G-400	6,000
	Summit X440	16
	Summit X460, X460-G2	6,000
	Summit X670, X480 (VIM3-40G4X), X670, X670-G2, X770	8,000
	Summit X480	245,760 (up to) <sup>b</sup>
<b>IPv6 routes with a mask greater than 64 bits in hardware</b> —maximum number of such IPv6 LPM routes in hardware.	BlackDiamond 8000 c-, e-, xm-series	256
	BlackDiamond 8000 xl-series	245,760 (up to) <sup>b</sup>
	BlackDiamond X8 series	256
	E4G-200, E4G-400	256
	Summit X440, X460, X460-G2, X670, X670-G2, X770, X480 (VIM3-40G4X)	256
	Summit X480	245,760 (up to) <sup>b</sup>

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IPv6 route sharing in hardware</b> —route mask lengths for which ECMP is supported in hardware.	Summit X460, X480, X670, X670V-48t	0-128
	E4G-200, E4G-400	0-128
	BlackDiamond 8800 (all I/O modules, except G48Te2), BlackDiamond X8 10G and 40G	0-128
	Summit X460-G2, X670-G2, X770	0-64 (> 64 single path only)
	BlackDiamond 8800 G48Te2, BlackDiamond X8 100G	0-64 (> 64 single path only)
	Summit X440	Single path only
<b>IPv6 routes in software</b> —maximum number of IPv6 routes in software.	BlackDiamond 8900 xl-series with 8900-MSM128 or MSM-48c	245,760 (up to) <sup>b</sup>
	All other BlackDiamond 8000 series hardware	25,000
	BlackDiamond X8 series	25,000
	Summit X460, X460-G2, X670, X670-G2, X770, E4G-200, E4G-400	25,000
	Summit X480	245,760 (up to) <sup>b</sup>
	Summit X440	256
<b>IP router interfaces</b> —maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	Summit X460-G2, X670, X770, X670-G2, and BlackDiamond X8	2,048
	Summit X440	254
	Summit X480, X460	512
	E4G-200, E4G-400	512
<b>IP multi-cast static routes</b> —maximum number of permanent multi-cast IP routes.	All platforms (except Summit X430, X440)	1,024
	Summit X430, X440	32
<b>IP unicast static routes</b> —maximum number of permanent IP unicast routes.	All platforms (except Summit X430, X440)	1,024
	Summit X430, X440	32



Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p><b>IP route sharing (maximum gateways)</b>—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.</p>	<p>All platforms, except Summit X430, X670, and BlackDiamond X8</p> <p>Summit X670, BlackDiamond X8</p> <p>Summit X430</p>	<p>2, 4, 8, 16, or 32</p> <p>2, 4, 6, 8, 16, 32, or 64</p> <p>N/A</p>
<p><b>IP route sharing (total destinations)</b>—maximum number of unique destinations used by multipath OSPF, OSPFv3, BGP, IS-IS, or static routes.</p> <p><b>NOTE:</b> For platforms with limit of 524,256, the total number of "destination+gateway" pairs is limited to 1,048,512. 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>	<p>BlackDiamond 8800 c-series</p> <p>BlackDiamond 8000 e-series</p> <p>BlackDiamond 8900 xl-series</p> <p>BlackDiamond 8900-40G6X-xm</p> <p>BlackDiamond X8 series</p> <p>E4G-200, E4G-400</p> <p>Summit X480</p> <p>Summit X670, X670-G2, X770, X480 (VIM3-40G4X)</p> <p>Summit X460-G2, X460</p>	<p>12,256</p> <p>480</p> <p>524,256 (up to)<sup>b</sup></p> <p>16,352</p> <p>16,000</p> <p>12,256</p> <p>524,256 (up to)<sup>b</sup></p> <p>16,352</p> <p>12,256</p>

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IP route sharing (total combinations of gateway sets)</b> —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 gateways is 64	510 1,022 254 126 62 30
	Summit X460, X460-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
<b>IP multinetting (secondary IP addresses)</b> —maximum number of secondary IP addresses per VLAN.	BlackDiamond 8800	64
	BlackDiamond X8	64
	All Summits, except X440, X430	255
	Summit X440	32
<b>IS-IS adjacencies</b> —maximum number of supported IS-IS adjacencies.	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	BlackDiamond 8900 xl-series	255
	Summit X440, X460, X460-G2, X480, X670, X670-G2, X770	128
	E4G-200 E4G-400	256 128
<b>IS-IS ECMP</b> —maximum number of equal cost multipath for IS-IS.	All platforms, except Summit X440, X430	2, 4, or 8
<b>IS-IS interfaces</b> —maximum number of interfaces that can support IS-IS.	All platforms, except Summit X440, X430	255

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>IS-IS IPv6 L2 routes</b> —recommended maximum number of IS-IS Level 2 routes.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	40,000
	Summit X480	15,000
	Summit X460, X460-G2, X670, X670-G2, X770	10,000
	E4G-200, E4G-400	10,000
<b>IS-IS IPv6 L1 routes in an L1/L2 router</b> —recommended maximum number of IS-IS Level 1 routes in a L1/L2 router.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	15,000
	Summit X480	15,000
	Summit X460, X460-G2, X670, X670-G2, X770, E4G-400	10,000
<b>IS-IS IPv4/IPv6 L1 routes in an L1 router</b> —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	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond 8900 xl-series	60,000
	Summit X480	40,000
	Summit X460, X460-G2, X670, X670-G2, X770	20,000
	E4G-200, E4G-400	20,000
<b>IS-IS IPv4/IPv6 L2 routes in an L2 router</b> —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	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond 8900 xl-series	60,000
	Summit X480	40,000
	Summit X460, X460-G2, X670, X670-G2, X770	20,000
	E4G-200, E4G-400	20,000
<b>IS-IS IPv4/IPv6 L1 routes in an L1/L2 router</b> —recommended maximum number of IS-IS Level 1 routes in a Level 1/Level 2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond 8900 xl-series	20,000
	Summit X460, X460-G2, X480, X670, X670-G2, X770	20,000
	E4G-200, E4G-400	20,000
<b>Jumbo frames</b> —maximum size supported for jumbo frames, including the CRC.	All platforms	9,216
<b>L2 VPN: VCCV (pseudowire Virtual Circuit Connectivity Verification) VPNs per switch</b> —maximum number of VCCV enabled VPLS VPNs.	All platforms, except Summit X440, X430	16

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>L2 VPN: VPLS MAC addresses</b> —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
	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	
<b>L2 VPN: VPLS VPNs</b> —maximum number of VPLS virtual private networks per switch.	BlackDiamond 8900 xl-series	1,023
	BlackDiamond 8900-40G6x-xm	1,023
	BlackDiamond X8 series	1,023
	E4G-200, E4G-400	1,023
	Summit X460, X460-G2, X480, X670, X670V-48t, X480 (40G VIM), X770, X670-G2	1,023
<b>L2 VPN: VPLS peers</b> —maximum number of VPLS peers per VPLS instance.	BlackDiamond 8900 xl-series	64
	BlackDiamond 8900-40G6x-xm	64
	BlackDiamond X8 series	64
	Summit X770, X670-G2, X670v-48t, X480, X460-G2	64
	Summit X670, X460	32
	E4G-200, E4G-400	32
<b>L2 VPN: LDP pseudowires</b> —maximum number of pseudowires per switch.	BlackDiamond 8900 xl-series	7,000
	BlackDiamond 8900-40G6X-xm	3,000
	BlackDiamond X8 series	7,000
	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	7,116
	Summit X460	1,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>L2 VPN: static pseudowires</b> — maximum number of static pseudowires per switch.	BlackDiamond 8900 xl-series, BlackDiamond X8	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	2,764
	E4G-400	6,860
<b>L2 VPN: Virtual Private Wire Service (VPWS) VPNs</b> — maximum number of virtual private networks per switch.	Summit X460	1,000
	Summit X480, X770	4,000
	Summit X480-40G VIM	2,047
	Summit 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	1,023
E4G-200, E4G-400	1,000	
<b>Layer-2 IPMC forwarding caches</b> — (IGMP/MLD/PIM snooping) in mac-vlan mode. <b>NOTE:</b> 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	15,000
	Summit X440	4,000
	Summit X770, X670-G2	77,500 <sup>i</sup>
	Summit X460-G2	24,576 <sup>k</sup>

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Layer-2 IPMC forwarding caches</b> — (IGMP/MLD/PIM snooping) in mixed-mode. <b>NOTE:</b> 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	15,000
	E4G-200 and E4G-400 cell site routers, Summit X460	8,000
	Summit X440	4,000
	Summit X770, X670-G2, X670	15,000 <sup>i</sup>
	Summit X460-G2	15,000 <sup>k</sup>
	Summit X480	8,000
<b>Layer-3 IPMC forwarding caches</b> — (PIM, MVR, PVLAN) in mixed- mode. <sup>9</sup> <b>NOTE:</b> 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 modules	64,000
	E4G-200 cell site routers, Summit X670	3,000
	E4G-400 cell site routers, Summit X460, X480, X670-48t	6,000
	Summit X440	192
	Summit X770, X670-G2	77,500 <sup>i</sup>
	Summit X460-G2	21,000 <sup>i</sup>

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Load sharing</b> —maximum number of loadsharing groups. <b>NOTE:</b> 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 with 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	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 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Load sharing</b> —maximum number of ports per load-sharing group. <b>NOTE:</b> * For custom algorithm ** For L2 and L3 algorithms <b>NOTE:</b> For a mix of Summit X770 and Summit X670 series switches in a stack, the limits are the Summit X670 limits.	BlackDiamond X8 series	64
	Summit X460-G2 (standalone)	32
	Summit X670 (standalone)	32 * 16 **
	Summit X670 (stacked)	64 *
	Summit X670-G2 (stacked)	16 **
	Summit X770 (standalone)	32
	Summit X670-G2 (standalone)	
	Summit X460-G2 (standalone)	
	Summit X770 (stacked) Summit X670-G2 (stacked) Summit X460-G2 (stacked)	64
	All other Summit series, SummitStacks, E4G cell site routers, and BlackDiamond 8000 series switches	8
<b>Logged messages</b> —maximum number of messages logged locally on the system.	All platforms	20,000
<b>MAC address learning rate</b> —hardware learning rate	E4G-200	22 msec
<b>MAC-based security</b> —maximum number of MAC-based security policies.	All platforms	1,024
<b>Mirroring (filters)</b> —maximum number of mirroring filters. <b>NOTE:</b> This is the number of filters across all the active mirroring instances.	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	All Summit series	128
	E4G cell site routers	128
<b>Mirroring (monitor port)</b> —maximum number of monitor ports.	All platforms	1
<b>Mirroring, one-to-many (filters)</b> —maximum number of one-to-many mirroring filters. <b>NOTE:</b> This is the number of filters across all the active mirroring instances.	BlackDiamond 8000 series	128
	BlackDiamond X8 series	128
	All Summit series	128
	E4G cell site routers	128
<b>Mirroring, one-to-many (monitor port)</b> —maximum number of one-to-many monitor ports.	All platforms	16

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Maximum mirroring instances</b> <b>NOTE:</b> 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	All platforms	16 (including default mirroring instance)
<b>MLAG ports</b> —maximum number of MLAG ports allowed.	BlackDiamond 8000 series BlackDiamond X8 series All Summit series, except X430 E4G cell site routers	768 768 768 768
<b>MLAG peers</b> —maximum number of MLAG peers allowed.	All platforms, except Summit X430	2
<b>MPLS RSVP-TE interfaces</b> —maximum number of interfaces.	All platforms	32
<b>MPLS RSVP-TE ingress LSPs</b> —maximum number of ingress LSPs.	All platforms	2,000
<b>MPLS RSVP-TE egress LSPs</b> —maximum number of egress LSPs.	All platforms	2,000
<b>MPLS RSVP-TE transit LSPs</b> —maximum number of transit LSPs.	All platforms	2,000
<b>MPLS RSVP-TE paths</b> —maximum number of paths.	All platforms, except Summit X670-G2 Summit X670-G2	1,000 2,000
<b>MPLS RSVP-TE profiles</b> —maximum number of profiles.	All platforms, except Summit X670-G2 Summit X670-G2	1,000 2,000
<b>MPLS RSVP-TE EROs</b> —maximum number of EROs per path.	All platforms	64

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>MPLS RSVP-TE fast reroute</b> —MPLS RSVP-TE fast reroute (FRR) switching time.	E4G-200	50 msec
<b>MPLS LDP peers</b> —maximum number of MPLS LDP peers per switch.	Summit X460, Summit X670 Summit X480, Summit X480 (40G VIM), X670V-48t, X770, X670v-48t BlackDiamond 8900 xl-series BlackDiamond 8900-40G6x-xm BlackDiamond X8 series Summit X670-G2, X460-G2 E4G-400, E4G-200	32 64 64 64 64 128 32
<b>MPLS LDP adjacencies</b> —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
<b>MPLS LDP ingress LSPs</b> —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 15,308 7116
<b>MPLS LDP-enabled interfaces</b> —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 2: Supported Limits (Continued)

Metric	Product	Limit
<b>MPLS LDP Sessions</b> —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
<b>MPLS LDP transit LSPs</b> —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	4,000
	Summit X670, X480 (VIM3-40G4x)	3,000
<b>MPLS LDP egress LSPs</b> —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	8,000
	Summit X670-G2, X460-G2	4,000
<b>MPLS static egress LSPs</b> —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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>MPLS static ingress LSPs</b> — 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
<b>MPLS static transit LSPs</b> — 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
<b>MSDP active peers</b> —maximum number of active MSDP peers.	BlackDiamond 8000 series	32
	BlackDiamond X8 series	64
	BlackDiamond 8900 series	64
	Summit X460, X480, X670, E4G- 400, X670-G2, X460-G2	16
	Summit X770	64
<b>MSDP SA cache entries</b> — maximum number of entries in SA cache.	BlackDiamond 8000 series	16,000
	BlackDiamond X8 series	16,000
	BlackDiamond 8900 series	16,000
	Summit X460, X480, X670, E4G- 400	8,000
	Summit X770	16,000
Summit X670-G2, X460-G2	14,000	
<b>MSDP maximum mesh groups</b> — maximum number of MSDP mesh groups.	BlackDiamond 8000 series	8
	BlackDiamond X8 series	16
	BlackDiamond 8900 series	16
	Summit X460, X480, X670, E4G- 400, X460-G2	4
	Summit X770, X670-G2	16

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Multi-cast listener discovery (MLD) IPv6 multi-cast data sender</b> —maximum number of IPv6 multi-cast streams supported on a switch. <sup>l 9</sup> <b>NOTE:</b> Assumes source-group-vlan mode. For additional limits, see: <ul style="list-style-type: none"> <li>• <b>Layer-2 IPMC forwarding caches</b>—(IGMP/MLD/PIM snooping) in mac-vlan mode. on page 46</li> <li>• <b>Layer-2 IPMC forwarding caches</b>— (IGMP/MLD/PIM snooping) in mixed-mode. on page 47</li> </ul>	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 series	3,000
	E4G-200	1,500
	E4G-400	3,000
	Summit X440	90
	Summit X460	3,000
	Summit X480	3,000
	Summit X670	1,500
Summit X770	4,096	
<b>Multi-cast listener discovery (MLD) snooping per-VLAN filters</b> —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 series	500
	E4G-400, Summit X460, X460-G2	1,000
	Summit X480	2,000
	Summit X440	250
	Summit X670, E4G-200	500
	Summit X770	1,200
Summit X670-G2	826	
<b>Multi-cast listener discovery (MLD)v1 subscribers</b> —maximum number of MLDv1 subscribers per port <sup>m</sup>	BlackDiamond 8800 c-series	500
	BlackDiamond xl-series	1,500
	BlackDiamond X8 Series	1,500
	Summit X440	750
	Summit X460, X460-G2, X480, X670, E4G-400	1,500
	Summit X770, X670-G2	4,000
<b>Multi-cast listener discovery (MLD)v1 subscribers</b> —maximum number of MLDv1 subscribers per switch <sup>m</sup>	BlackDiamond 8800 series	10,000
	BlackDiamond X8 series	10,000
	Summit X440	5,000
	Summit X460, X480, X670, E4G-400, X460-G2	10,000
	Summit X770, X670-G2	30,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Multi-cast listener discovery (MLD)v2 subscribers</b> —maximum number of MLDv2 subscribers per port <sup>m</sup>	BlackDiamond 8800 c-series	500
	BlackDiamond xl series	2,500
	BlackDiamond X8 series	2,000
	Summit X440, SummitStack	1,000
	Summit X460, X480, X670, E4G-400, X460-G2	2,000
	Summit X770, X670-G2	4,000
<b>Multi-cast listener discovery (MLD)v2 subscribers</b> —maximum number of MLDv2 subscribers per switch <sup>m</sup>	BlackDiamond 8800 series	10,000
	BlackDiamond xl series	10,000
	Summit X440, SummitStack	5,000
	Summit X460, X480, X670, E4G-400, X460-G2	10,000
	Summit X770, X670-G2	30,000
<b>Multi-cast listener discovery (MLD)v2 maximum source per group</b> —maximum number of source addresses per group	All platforms, except Summit X430	200
<b>Multi-cast VLAN registration (MVR)</b> —maximum number of MVR senders per switch (IP multi-cast compression disabled). <b>NOTE:</b> Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 <sup>f</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 <sup>f</sup>
	8900-G96T-c modules	4,096 <sup>f</sup>
	8900 xl-series	4,096 <sup>f</sup>
	8900-40G6X-xm	3,000 <sup>g</sup>
	BlackDiamond X8 a-series modules	4,096
	BlackDiamond X8-100G4X modules	16,000
	E4G-200	2,048
	E4G-400	500 <sup>g</sup>
	Summit X440	1,024
	Summit X480	2,048
	Summit X460	2,048
	Summit X670	
VIM4-40G4x	3,000 <sup>g</sup>	
Summit X770, X670-G2	4,096	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Multi-cast VLAN registration (MVR)</b> —maximum number of MVR senders per switch (IP multi-cast compression enabled). <b>NOTE:</b> Assumes source-group-vlan mode. For additional limits, see <a href="#">Layer-3 IPMC forwarding caches</a> —(PIM, MVR, PVLAN) in mixed-mode.9 on page 47.	BlackDiamond 8800 c-series	6,000 <sup>g</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 c-series	6,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	12,000 <sup>b</sup>
	BlackDiamond X8 a-series modules	6,000 <sup>g</sup>
	BlackDiamond X8-100G4X modules	64,000
	8900-40G6X-xm module	3,000 <sup>g</sup>
	Summit X440	192 <sup>g</sup>
	Summit X460, E4G-400	6,000 <sup>g</sup>
	Summit X480	12,000 <sup>b</sup>
	Summit X670	
	VIM4-40G4x	3,000 <sup>g</sup>
Summit X770	6,300	
Summit X670-G2	10,000	
<b>Network login</b> —maximum number of clients being authenticated on MAC-based VLAN enabled ports.	BlackDiamond 8000 series (clients per module/per system)	1,024
	BlackDiamond X8 series	1,024
	Summit series	1,024
<b>Network login</b> —maximum number of dynamic VLANs.	All platforms	2,000
<b>Network login VLAN VSAs</b> —maximum number of VLANs a client can be authenticated on at any given time.	All platforms	10
<b>OSPFv2/v3 ECMP</b> —maximum number of equal cost multipath OSPFv2 and OSPFv3.	All platforms, except Summit X440, X430, and E4G-200)	16
	E4G-200	8
<b>OSPFv2 areas</b> —as an ABR, how many OSPF areas are supported within the same switch.	All platforms (except X430, X440)	8
<b>OSPFv2 external routes</b> —recommended maximum number of external routes contained in an OSPF LSDB.	BlackDiamond 8000 series	20,000
	BlackDiamond 8900 xl-series	130,000
	BlackDiamond X8 series	20,000
	Summit X460, X670, X770, X670-G2, X460-G2	5,000
	Summit X480	130,000
	E4G-400	5,000
	E4G-200	5,000



Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>OSPFv2 inter- or intra-area routes</b> —recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	BlackDiamond 8000 series	7,000
	BlackDiamond 8900 xl-series	7,000
	BlackDiamond X8 series	7,000
	Summit X460, X670, X670-G2, X460-G2	2,000
	E4G-400	2,000
	Summit X480, X770	7,000
<b>OSPFv2 interfaces</b> —recommended maximum number of OSPF interfaces on a switch.	<b>NOTE:</b> Active interfaces limit, with Advanced Edge license. (See below for Core license limits.) All platforms (except X430)	4
	All platforms (except X430 and X440) with Core license or higher (active interfaces only)	400
<b>OSPFv2 links</b> —maximum number of links in the router LSA.	All platforms, except Summit X770 and X430	400
	Summit X770	419
<b>OSPFv2 neighbors</b> —maximum number of supported OSPF adjacencies.	BlackDiamond 8000 series	128
	BlackDiamond 8900 xl-series	255
	BlackDiamond X8 Series	255
	Summit X460, X670, X770, X440, X670-G2, X460-G2	128
	Summit X480	255
	E4G-400, E4G-200	128
<b>OSPFv2 routers in a single area</b> —recommended maximum number of routers in a single OSPF area.	BlackDiamond 8000 series	100
	BlackDiamond 8900 xl-series	200
	BlackDiamond X8 series	100
	Summit X460, X670, X770, X670-G2, X460-G2	50
	Summit X480	200
	E4G-400	50
<b>OSPFv2 virtual links</b> —maximum number of supported OSPF virtual links.	All platforms (except X430 and X440) with Core license or higher	32
<b>OSPFv3 areas</b> —as an ABR, the maximum number of supported OSPFv3 areas.	All platforms (except X430 and X440) with Core license or higher	16
<b>OSPFv3 external routes</b> —recommended maximum number of external routes.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond 8900 xl-series	60,000
	Summit X460, X670, X770, X670-G2, X460-G2	10,000
	Summit X480	60,000
	E4G-400	10,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>OSPFv3 inter- or intra-area routes</b> —recommended maximum number of inter- or intra-area routes.	BlackDiamond 8000 series	6,000
	BlackDiamond X8 series	6,000
	BlackDiamond 8900 xl-series	6,000
	Summit X460, X670, X770, X670-G2, X460-G2	3,000
	Summit X480	6,000
	E4G-400	3,000
<b>OSPFv3 interfaces</b> —maximum number of OSPFv3 interfaces.	<b>NOTE:</b> Active interfaces only, with Advanced Edge license. (See below for Core license limits.) All platforms (except X430)	4
	<b>NOTE:</b> With Core license or higher. (See above for Advanced Edge license limits.) BlackDiamond 8000 series BlackDiamond X8 series BlackDiamond 8900 xl-series Summit X460, X670, X770 Summit X480 Summit X670-G2, X460-G2 E4G-200, E4G-400	256 256 384 128 384 256 256
<b>OSPFv3 neighbors</b> —maximum number of OSPFv3 neighbors.	BlackDiamond 8000 series	64
	BlackDiamond X8 series	64
	BlackDiamond 8900 xl-series	128
	Summit X460, X670, X770, X670-G2, X460-G2	64
	Summit X480	128
	E4G-400	64
<b>OSPFv3 virtual links</b> —maximum number of OSPFv3 virtual links supported.	All platforms (except X430 and X440) with Core license or higher	16

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p><b>PIM IPv4 snooping</b>—maximum number of (S,G) entries programmed in the hardware (IP multi-cast compression disabled).</p> <p><b>NOTE:</b> Assumes source-group-vlan mode.</p>	BlackDiamond 8800 c-series	2,048 <sup>f</sup>
	BlackDiamond 8000 e-series	500 <sup>f</sup>
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 <sup>f</sup>
	8900-G96T-c modules	4,096 <sup>f</sup>
	8900 xl-series	4,096 <sup>f</sup>
	8900-40G6X-xm	3,000 <sup>g</sup>
	BlackDiamond X8 a-series modules	4,096
	BlackDiamond X8-100G4X modules	16,384
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64
	Summit X460, X460-G2	2,048
	Summit X480	4,096
Summit X670	3,000 <sup>g</sup>	
Summit X670-G2	3,000	
Summit X770	4,096	
<p><b>PIM IPv4 snooping</b>—maximum number of (S,G) entries programmed in the hardware (IP multi-cast compression enabled).</p> <p><b>NOTE:</b> Assumes source-group-vlan mode.</p> <p>For additional limits, see:</p> <ul style="list-style-type: none"> <li>• <b>Layer-2 IPMC forwarding caches</b>—(IGMP/MLD/PIM snooping) in mac-vlan mode. on page 46</li> <li>• <b>Layer-2 IPMC forwarding caches</b>—(IGMP/MLD/PIM snooping) in mixed-mode. on page 47</li> </ul>	BlackDiamond 8800 c-series	6,000 <sup>g</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 c-series	6,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	12,000 <sup>b</sup>
	BlackDiamond X8 a-series modules	6,000 <sup>g</sup>
	BlackDiamond X8-100G4X modules	64,000 <sup>g</sup>
	E4G-200	3,000 <sup>g</sup>
	E4G-400	6,000 <sup>g</sup>
	8900-40G6X-xm	3,000 <sup>g</sup>
	Summit X440	192 <sup>g</sup>
	Summit X480	12,000 <sup>b</sup>
	Summit X460, X670-G2, X460-G2	6,000 <sup>g</sup>
	Summit X670	3,000 <sup>g</sup>
	Summit X770	66,500

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>PIM IPv4—maximum routes—</b> maximum number of (S,G) entries installed in the hardware (IP multi-cast compression disabled). <b>NOTE:</b> Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 <sup>f</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 <sup>f</sup>
	8900-G96T-c modules	4,096 <sup>f</sup>
	8900 xl-series	4,096 <sup>f</sup>
	8900-40G6X-xm	3,000 <sup>g</sup>
	BlackDiamond X8 a-series modules	4,094
	BlackDiamond X8-100G4X modules	16,384
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64 <sup>g</sup>
	Summit X480, X670-G2, X460-G2	4,096
Summit X460	2,048	
Summit X670	3,000 <sup>g</sup>	
Summit X770	4,096	
<b>PIM IPv4—maximum routes—</b> maximum number of (S,G) entries installed in the hardware (IP multi-cast compression enabled). <b>NOTE:</b> Assumes source-group-vlan mode. For additional limits, see <b>Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode.</b> <sup>9</sup> on page 47.	BlackDiamond 8800 c-series	6,000 <sup>g</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 c-series	6,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	12,000 <sup>b</sup>
	BlackDiamond X8 a-series modules	6,000 <sup>f</sup>
	BlackDiamond X8-100G4X	64,000 <sup>f</sup>
	E4G-200	3,000 <sup>g</sup>
	E4G-400	6,000 <sup>g</sup>
	8900-40G6X-xm modules	3,000 <sup>g</sup>
	Summit X440	192
	Summit X480	12,000 <sup>b</sup>
	Summit X460	6,000 <sup>g</sup>
	Summit X670	3,000 <sup>g</sup>
	Summit X770, X670-G2	66,500
	Summit X460-G2	21,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>PIM IPv4-SSM (maximum SSM routes)</b> —maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multi-cast compression disabled). <b>NOTE:</b> Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 <sup>f</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 <sup>f</sup>
	8900-G96T-c modules	4,096 <sup>f</sup>
	8900 xl-series	15,000
	8900-40G6X-xm	3,000 <sup>g</sup>
	BlackDiamond X8 a-series modules	4,094
	BlackDiamond X8-100G4X modules	16,384
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64
	Summit X480, X670-G2, X460-G2	4,096
Summit X460	2,048	
Summit X670	3,000 <sup>g</sup>	
Summit X770	4,096	
<b>PIM IPv4-SSM (maximum SSM routes)</b> —maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multi-cast compression enabled). <b>NOTE:</b> Assumes source-group-vlan mode. For additional limits, see: <b>Layer-3 IPMC forwarding caches</b> —(PIM, MVR, PVLAN) in mixed-mode. <sup>g</sup> on page 47	BlackDiamond 8800 c-series	6,000 <sup>g</sup>
	BlackDiamond 8000 e-series	500 <sup>g</sup>
	BlackDiamond 8900 c-series	6,000 <sup>g</sup>
	BlackDiamond 8900 xl-series	12,000 <sup>b</sup>
	BlackDiamond X8 a-series modules	6,000 <sup>g</sup>
	BlackDiamond X8-100G4X modules	64,000 <sup>g</sup>
	E4G-200	3,000 <sup>g</sup>
	E4G-400	6,000 <sup>g</sup>
	8900-40G6X-xm	3,000 <sup>g</sup>
	Summit X440	192 <sup>g</sup>
	Summit X480	12,000 <sup>b</sup>
	Summit X460	6,000 <sup>g</sup>
	Summit X670	3,000 <sup>g</sup>
	Summit X770	66,500
	Summit X670-G2	60,000
Summit X460-G2	21,000	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>PIM IPv6 (maximum routes)</b> —maximum number of (S,G) entries installed in the hardware. <b>NOTE:</b> 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 modules	3,000
	BlackDiamond X8-100G4X modules	64,000 <sup>d</sup>
	E4G-200	1,500
	E4G-400	3,000
	Summit X440	90
	Summit X460, X460-G2, X480	3,000
	Summit X670	1,500
Summit X770, X670-G2	38,000	
<b>PIM IPv4 (maximum interfaces)</b> —maximum number of PIM active interfaces.	All platforms, except Summit X430, X440	512
<b>PIM IPv4 (maximum interfaces)</b> —maximum number of PIM snooping enabled interfaces.	All platforms, except Summit X430, X440	512
<b>PIM IPv4 Limits</b> —maximum number of multi-cast groups per rendezvous point	All platforms, except Summit X430, X440	180
<b>PIM IPv4 Limits</b> —maximum number of multi-cast sources per group	All platforms, except Summit X430, X440	175
<b>PIM IPv4 Limits</b> —maximum number of dynamic rendezvous points per multi-cast group	All platforms, except Summit X430, X440	145
<b>PIM IPv4 Limits</b> —static rendezvous points	All platforms, except Summit X430, X440	32
<b>PIM IPv6 (maximum interfaces)</b> —maximum number of PIM active interfaces	All platforms, except Summit X430, X440	512
<b>PIM IPv6 Limits</b> —maximum number of multicast group per rendezvous point	All platforms, except Summit X430, X440	70
<b>PIM IPv6 Limits</b> —maximum number of multicast sources per group	All platforms, except Summit X430, X440	43
<b>PIM IPv6 Limits</b> —maximum number of dynamic rendezvous points per multicast group	All platforms, except Summit X430, X440	64
<b>PIM IPv6 Limits</b> —maximum number of secondary address per interface	All platforms, except Summit X430, X440	70

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>PIM IPv6 Limits</b> —static rendezvous points	BlackDiamond X8 series	32
	BlackDiamond 8800 series	32
	BlackDiamond 8900 series	32
	Summit X770, X670-G2, X670v-48t	32
	Summit X670, X480, X460-G2, X460	70
	E4G-200, E4G-400	70
<b>Policy-based routing (PBR) redundancy</b> —maximum number of flow-redirects.	All platforms	256 <sup>n</sup>
<b>Policy-based routing (PBR) redundancy</b> —maximum number of next hops per each flow-direct.	All platforms	32 <sup>n</sup>
<b>Port-specific VLAN tags</b> —maximum number of port-specific VLAN tags	All platforms	1,023
<b>Port-specific VLAN tags</b> —maximum number of port-specific VLAN tag ports	BlackDiamond X8 and BlackDiamond 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	
<b>Private VLANs</b> —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	767
	Summit X770	103
	Summit X670-G2, X670v-48t	63
	Summit X670	47
	Summit X480	23
	Summit X460-G2, X460	53
	Summit X440	25
	Summit X430	27
	E4G-200	11
	E4G-400	33

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Private VLANs</b> —maximum number of private VLANs with an IP address on the network VLAN.  <b>NOTE:</b> 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, X460, X480, X670-G2, X460-G2	1,024
	Summit X670, X480, X460	512
	Summit X440	256
	Summit X440	127
	E4G-200, E4G-400	512
<b>Private VLANs</b> —maximum number of private VLANs in an L2-only environment.	BlackDiamond 8800 c-, e-series	384
	BlackDiamond 8900 series	2,046
	BlackDiamond X8 series	2,046
	E4G-200	597
	E4G-400	1,280
	Summit X440	127
	Summit X480	2,046
	Summit X670	597
	Summit X460	820
Summit X770, X670-G2, X460-G2	1,280	
<b>PTP/1588v2 Clock Ports</b>	Summit X770, X460-G2, X670-G2, and E4G-200, E4G-400 cell site routers	32 for boundary clock
		1 for ordinary clock
<b>PTP/1588v2 Clock Instances</b>	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	2 combinations: <ul style="list-style-type: none"> <li>• Transparent clock + ordinary clock</li> <li>• Transparent clock + boundary clock</li> </ul>
<b>PTP/1588v2 Unicast Static Slaves</b>	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	40 entries per clock port
<b>PTP/1588v2 Unicast Static Masters</b>	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	10 entries per clock type
<b>Route policies</b> —suggested maximum number of lines in a route policy file.	All platforms	10,000
<b>RIP Learned Routes</b> —maximum number of RIP routes supported without aggregation.	All platforms, except Summit X430	10,000
<b>RIP neighbors</b> —maximum number of RIP neighbors.	E4G-200	256



Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>RIP interfaces on a single router</b> —recommended maximum number of RIP routed interfaces on a switch.	BlackDiamond 8000 series	256
	BlackDiamond X8 series	256
	BlackDiamond 8900 xl-series	384
	Summit X440	128
	Summit X460, X670-G2, X460-G2	256
	Summit X480	384
	Summit X670, X770	256
E4G-400	256	
<b>RIPng learned routes</b> —maximum number of RIPng routes.	BlackDiamond 8000 series	3,000
	BlackDiamond X8 series	3,000
	BlackDiamond 8900 xl-series	5,000
	Summit X480	5,000
	Summit X460, X670, X670-G2, X460-G2, X770	3,000
	E4G-200	3,000
<b>Spanning Tree (maximum STPDs)</b> —maximum number of Spanning Tree Domains on port mode EMISTP.	All platforms (except Summit X430 and Summit X440)	64
	Summit X440	32
	Summit X430	16
<b>Spanning Tree PVST+</b> —maximum number of port mode PVST domains. <b>NOTE:</b> <ul style="list-style-type: none"> <li>Maximum of 10 active ports per PVST domain when 256 PVST domains are configured.</li> <li>Maximum of 7 active ports per PVST domain when 128 PVST domains are configured.</li> </ul>	BlackDiamond X8 and 8900 series switches	256
	Summit X670, X770, X670-G2	256
	Summit X460, X480, X440, X460-G2	128
	Summit X430	50
	E4G-400	128
<b>Spanning Tree</b> —maximum number of multiple spanning tree instances (MSTI) domains.	All platforms (except Summit X430 and Summit X440)	64
	Summit X440	32
	Summit X430	5
<b>Spanning Tree</b> —maximum number of VLANs per MSTI. <b>NOTE:</b> Maximum number of 10 active ports per VLAN when all 500 VLANs are in one MSTI.	BlackDiamond X8	500
	BlackDiamond 8800	500
	BlackDiamond 8900 MSM 128/XL	500
	Summit X770, X670-G2, X670v-48t, X670	500
	Summit X480, X460-G2, X460	600
	E4G-200	500
	E4G-400	600
	Summit X440	250
	Summit X430	100

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>Spanning Tree</b> —maximum number of VLANs on all MSTP instances.	BlackDiamond X8	1,000
	BlackDiamond 8800	1,000
	BlackDiamond 8900 MSM 128/XL	1,000
	Summit X770	1,024
	Summit X670-G2, X670v-48t, X670, X480	1,000
	Summit X460-G2, X460	1,024
	E4G-200	1,000
	E4G-400	1,024
	Summit X440	500
Summit X430	200	
<b>Spanning Tree (802.1d domains)</b> —maximum number of 802.1d domains per port.	All platforms	1
<b>Spanning Tree (number of ports)</b> —maximum number of ports including all Spanning Tree domains.	All platforms (except Summit X430 and Summit X440)	4,096
	Summit X440	2,048
	Summit X430	1,024
<b>Spanning Tree (maximum VLANs)</b> —maximum number of STP-protected VLANs (dot1d and dot1w).	BlackDiamond X8	1,024
	BlackDiamond 8800	1,024
	BlackDiamond 8900 MSM 128/XL	1,024
	Summit X770	1,024
	Summit X670-G2, X670v-48t, X670, X480	560
	Summit X460-G2, X460	600
	E4G-200	500
	E4G-400	600
	Summit X440	500
Summit X430	128	
<b>SSH (number of sessions)</b> —maximum number of simultaneous SSH sessions.	All platforms	8
<b>Static MAC multi-cast FDB entries</b> —maximum number of permanent multi-cast MAC entries configured into the FDB.	BlackDiamond 8000 c-, e-, xl-series	1,024
	BlackDiamond X8 series	1,024
	All Summits	1,024
	E4G-200, E4G-400	1,024
<b>Syslog servers</b> —maximum number of simultaneous syslog servers that are supported.	All platforms	4
<b>Telnet (number of sessions)</b> —maximum number of simultaneous Telnet sessions.	All platforms	8

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>VLANs</b> —includes all VLANs. <b>NOTE:</b> 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
<b>VLANs</b> —maximum number of port-specific tag VLANs.	BlackDiamond 8800 xl-series only, BlackDiamond X8 series	1,023
	BlackDiamond X8 xl-series	4,093
	Summit X460, X770, X480, E4G-400, X670-G2, X460-G2	4,093
	Summit X670, X670V-48t	1,023
	E4G-400	4,093
	E4G-200	2,047
<b>VLANs</b> —maximum number of port-specific tag VLAN ports	Summit X460, X670, X670V-48t, X460-G2, BlackDiamond 8800 xl- series only, BlackDiamond X8, E4G- 400, E4G-200	4,096
	BlackDiamond X8 xl-series	32,767
	Summit X770, X670-G2	8,192
	Summit X480	16,383
<b>VLANs (Layer 2)</b> —maximum number of Layer 2 VLANs.	All platforms	4,094
<b>VLANs (Layer 3)</b> —maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	Summit X460-G2, X670, X770, X670-G2, and BlackDiamond X8	2,048
	Summit X440	254
	Summit X480, X460	512
	E4G-200, E4G-400	512
<b>VLANs (maximum active port-based)</b> —maximum active ports per VLAN when 4,094 VLANs are configured with default license.	BlackDiamond X8	32
	BlackDiamond 8800 series	32
	Summit X770, X670-G2, X670v-48t, X670, X480, X460-G2, X460	32
	E4G-200	12
	E4G-400	32
	Summit X440	13
Summit X430	2	
<b>VLANs (maximum active protocol-sensitive filters)</b> —number of simultaneously active protocol filters in the switch.	All platforms	15

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>VLAN translation</b> —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 (383) 8900-G96T-c modules (767)
	Summit X770	103
	Summit X670-G2, X670v-48t	63
	Summit X670	47
	Summit X480	23
	Summit X460-G2	53
	Summit X460	57
	E4G-200	11
	E4G-400	33
	Summit X440	25
Summit X430	27	
<b>VLAN translation</b> —maximum number of translation VLAN pairs with an IP address on the translation VLAN. <b>NOTE:</b> 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	1,024
	Summit X670, X480, X460	512
	Summit X460-G2	1,024
	E4G-200, E4G-400	512
	Summit X440	127
<b>VLAN translation</b> —maximum number of translation VLAN pairs in an L2-only environment.	BlackDiamond 8800 c-, e-series	384
	BlackDiamond 8900 xl-series	2,046
	BlackDiamond X8 series	2,046
	Summit X460, X670-G2, X460-G2	2,000
	E4G-400, E4G-200	2,000
	Summit X440	512
	Summit X480, X670, X770	2,046
	Summit X430	100
<b>VRRP (v2/v3-IPv4) (maximum instances)</b> —maximum number of VRRP instances for a single switch, with Advanced Edge license or higher.	BlackDiamond X8, 8800 c-series MSM-48c, and BlackDiamond 8900 xl-series 8900-MSM128	511
	Summit X770, X670, X670-G2, X460-G2, X480	511
	E4G-200, E4G-400	128
	Summit X460	255
	Summit X440	32

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>VRRP (v3-IPv6) (maximum instances)</b> —maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. (VRRP-VRRPv3-IPv6)	BlackDiamond X8, 8800 c-series MSM-48c, and BlackDiamond 8900 xl-series 8900-MSM128	511
	Summit X770, X670, X670-G2, X460-G2	511
	E4G-200, E4G-400	255
	Summit X460	255
	Summit X480 Summit X440	255 15
<b>VRRP (v2/v3-IPv4/IPv6) (maximum VRID)</b> —maximum number of unique VRID numbers per switch.	All platforms with Advanced Edge license or higher, except Summit X430	7
<b>VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN)</b> —maximum number of VRIDs per VLAN.	All platforms with Advanced Edge license or higher, except for Summit X430	7
<b>VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks)</b> —maximum number of ping tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
<b>VRRP (maximum ping tracks)</b> —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	1
	Summit X440 Hello interval: 1 second	1
<b>VRRP (v3-IPv6) (maximum ping tracks)</b> —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	1 (IPv6)
	Summit X440 Hello interval: 1 second	1 (IPv6)
<b>VRRP (v2/v3-IPv4/IPv6) (maximum iproute tracks)</b> —maximum number of IP route tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
<b>VRRP (v2/v3-IPv4/IPv6)</b> —maximum number of VLAN tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<b>XML requests</b> —maximum number of XML requests per second. <b>NOTE:</b> 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 8900 series with 100 DACLs with 500 DACLs	10 3
	Summit X480, X670 with 100 DACLs with 500 DACLs	4 1
<b>XNV authentication</b> —maximum number of VMs that can be processed (combination of local and network VMs).	All platforms, except Summit X430	2,048
<b>XNV database entries</b> —maximum number of VM database entries (combination of local and network VMs).	All platforms, except Summit X430	16,000
<b>XNV database entries</b> —maximum number of VPP database entries (combination of local and network VPPs).	All platforms, except Summit X430	2,048
<b>XNV dynamic VLAN</b> —Maximum number of dynamic VLANs created (from VPPs /local VMs).	All Platforms, except Summit X430	2,048
<b>XNV local VPPs</b> —maximum number of XNV local VPPs.	All platforms, except Summit X430 Ingress Egress	2,048 512
<b>XNV policies/dynamic ACLs</b> —maximum number of policies/dynamic ACLs that can be configured per VPP. <sup>o</sup>	All platforms, except Summit X430 Ingress Egress	8 4
<b>XNV network VPPs</b> —maximum number of XNV network VPPs. <sup>o</sup>	All platforms, except Summit X430 Ingress Egress	2,048 512

- The table shows the total available.
- Limit depends on setting configured for `configure forwarding external-tables`.
- When there are BFD sessions with minimal timer, sessions with default timer should not be used.
- Based on forwarding internal table configuration "more I2".
- Effective capacity varies based on actual MAC addresses and VLAN IDs used and hash algorithm selected.
- Applies only if all enabled BlackDiamond 8000 I/O modules are BlackDiamond 8000 c-, xl-, or xm-series modules.
- Effective capacity varies based on actual IP addresses and hash algorithm selected, but is higher for BlackDiamond 8000 c-, xl-, xm-series modules, BlackDiamond X8, E4G cell site routers, and Summit X460 and X480 switches compared to BlackDiamond 8000 e-series modules.
- For the MVR feature in the BlackDiamond X8 series switches, the number of senders applies only when there are few egress VLANs with subscribers. If there are many VLANs with subscribers, the limit is substantially less. Only 500 senders are supported for 100 VLANs. It is not recommended to exceed these limits.
- Based on forwarding internal table configuration "I3-and-ipmc".
- The limit depends on setting configured with `configure iproute reserved-entries`.
- Based on forwarding internal table configuration "I2-and-I3".

- l. The IPv4 and IPv6 multi-cast entries share the same hardware tables, so the effective number of IPv6 multi-cast entries depends on the number of IPv4 multi-cast entries present and vice-versa.
- m. If IGMP and MLD are simultaneously configured on the switch, the number of effective subscribers supported would be appropriately lessened.
- n. Sum total of all PBR next hops on all flow redirects should not exceed 1024.
- o. The number of XNV authentications supported based on system ACL limitations.



# 3 Open Issues, Known Behaviors, and Resolved Issues

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

This chapter contains the following sections:

- [Open Issues on page 73](#)
- [Known Behaviors on page 76](#)
- [Resolved Issues in ExtremeXOS 15.6.4-Patch1-5 on page 77](#)
- [Resolved Issues in ExtremeXOS 15.6.4-Patch1-3 on page 78](#)
- [Resolved Issues in ExtremeXOS 15.6.4 on page 81](#)
- [Resolved Issues in ExtremeXOS 15.6.3-Patch1-9 on page 83](#)
- [Resolved Issues in ExtremeXOS 15.6.3-Patch1-8 on page 85](#)
- [Resolved Issues in ExtremeXOS 15.6.3-Patch1-5 on page 87](#)
- [Resolved Issues in ExtremeXOS 15.6.3-Patch1-3 on page 89](#)
- [Resolved Issues in ExtremeXOS 15.6.3 on page 93](#)
- [Resolved Issues in ExtremeXOS 15.6.2-Patch1-6 on page 94](#)
- [Resolved Issues in ExtremeXOS 15.6.2-Patch1-1 on page 97](#)
- [Resolved Issues in ExtremeXOS 15.6.2 on page 99](#)
- [Resolved Issues in ExtremeXOS 15.6.1 on page 104](#)

## Open Issues

The following are new open issues for supported features found in ExtremeXOS 15.6.4-Patch1-5.

**Table 3: Open Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
xos0056469	If the reauth timer expires, the client becomes unauthenticated and then cannot access the Internet. Administrator intervention is needed to reinitialize the client ( <code>disable netlogin web-based</code> ) so that the client can log on again and connect to the Internet. The desired behavior should be that after becoming unauthenticated, you should be directed to the logon page to supply logon credentials and be granted network access again.
xos0057913	Switch stops responding and you cannot save or show a configuration after disabling BGP and deleting virtual routing and forwarding (VRF).

**Table 3: Open Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
xos0055511	<p>While configuring STP (802.1d) with port-encapsulation mode as EMISTP where the L2PT-enabled VMAN and access VLAN have the same tag, the designated bridge is not accepting the L2PT tunneled BPDUs from the root bridge, and thus causes a loop (designated bridge also becomes a root bridge).</p> <p>This problem does not occur:</p> <ul style="list-style-type: none"> <li>• When the access VLAN's tag and the L2PT-enabled VMAN's tag are different.</li> <li>• Without any L2PT configured, with the same tag used for the access VLAN and provider-edge VMAN.</li> <li>• When using Per-VLAN Spanning Tree Plus (PVST+), regardless of same or different tags.</li> </ul>
xos0055517	L2PT filter match criteria for tunneling does not work for destination-mac, etype, and offset combination when ingress traffic on access ports has a VLAN tag.
<b>BlackDiamond 8800 Series Switches</b>	
xos0056118	TCP red packets are not dropped when the wredprofile is set up to discard them.
<b>BlackDiamond X8 Series Switches</b>	
xos0055864	When you configure a large ACL (999 rules) on a port, slow path traffic rate drops down to 1,00–2,00 packets/second from a rate of 55,000–60,000 packets/second.
<b>Summit Family Switches</b>	
xos0057835	In SummitStacks, clear-flow sampling period is incorrectly calculated.
xos0057250	Restarting BGP processes or rebooting the switch re-enables BGP processes running on disabled VPN-VRFs.
<b>SummitStack</b>	
xos0056433	In a stack with Summit X770 series switches as the master and backup with broadcast traffic running from master/backup to standby slots, standby slots may go to failed state when a stack failover occurs.
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.
<b>Summit X430 Series Switches</b>	
xos0056063	Summit X430-48t switches can only create a maximum of 3,932 VLANs; Summit X430-24t switches can only create a maximum of 3,996 VLANs. The published limit for Summit X430 series switches is 4,094 VLANs. Other Summit X430 series switches can meet the 4,094 limit.

**Table 3: Open Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
xos0055518	<p>On Summit X460 series switches with OpenFlow enabled, switches boot up with the following error messages:</p> <pre> “Error while loading structure &lt;cfgTechSupport&gt;&lt;basic&gt;1&lt;/basic&gt;&lt;hour&gt;0&lt;/ hour&gt;&lt;srcipaddress&gt;10.68.63.86&lt;/ srcipaddress&gt;&lt;hostname&gt;&lt;![CDATA[12.38.14.200]]&gt;&lt;/ hostname&gt;&lt;sslEnabled&gt;0&lt;/sslEnabled&gt;&lt;port&gt;800&lt;/ port&gt;&lt;daily&gt;0&lt;/daily&gt;&lt;isDefault&gt;1&lt;/ isDefault&gt;&lt;bootup&gt;1&lt;/bootup&gt;&lt;automatic&gt;1&lt;/ automatic&gt;&lt;nameIndex&gt;&lt;![CDATA[12.38.14.200]]&gt;&lt;/ nameIndex&gt;&lt;errorDetected&gt;0&lt;/ errorDetected&gt;&lt;vrName&gt;&lt;![CDATA[VR-Mgmt]]&gt;&lt;/vrName&gt;&lt;/ cfgTechSupport&gt;: Source IP address 10.68.63.86 does not belong to the VR VR-Mgmt. &lt;cfgTechSupport&gt;”  “Error while loading structure &lt;openflowGlobalConfig&gt;&lt;numVlans&gt;2&lt;/ numVlans&gt;&lt;version&gt;0&lt;/version&gt;&lt;fdb&gt;0&lt;/ fdb&gt;&lt;isEnabled&gt;1&lt;/isEnabled&gt;&lt;/openflowGlobalConfig&gt;: Error: Invalid OpenFlow version&lt;openflowGlobalConfig&gt;” </pre>
<b>Summit X670 Series Switches</b>	
xos0059426	On Summit X670v (80G) series stacks, system may stop responding and display “process epm pid 1118 signal 6” error on master slot when running failover with mirroring configuration and multi-cast traffic.
xos0055917	With OpenFlow v1.3 enabled, the switch stops responding after executing the <code>disable openflow</code> command. This problem does not occur when the switch has OpenFlow v1.0 enabled.
<b>Summit X670-G2 Series Switches</b>	
xos0056400	BGP does not elect the best path after failover on SummitStacks.
<b>Summit X770 Series Switches</b>	
xos0056006	When broadcast traffic is flowing, after disabling, and then enabling a port, the port can take more than 15 seconds to become active.
<b>ERPS</b>	
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.
<b>OpenFlow</b>	
xos0056061	The OpenFlow switch is not actively sending out hello messages when the 3-way TCP connection is established.

## Known Behaviors

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

**Table 4: Known Issues, Platform-Specific, and Feature Change Requets (CRs)**

CR Number	Description
<b>General</b>	
xos0056975	PIM-SM: Traffic loss occurs for up to 60 seconds when the ingress port on a non-designated, first-hop router (FHR) is disabled.
xos0055977	After disabling the EAPS shared port and a segment port in a partner node, traffic fails to flow for 30 seconds after enabling this segment port.
xos0057187	BFD does not work if the client is a VRF leaked route in another VR or VRF. Configuring BFD session to a virtual IP is unsupported. <b>Workaround:</b> Configure a dummy static route in VR.
<b>BlackDiamond Series Switches</b>	
xos0055816	BlackDiamond 8800 c-series modules support IPv6 routes with masks greater than 64 only for virtual router "VR-Default". The following error appears if a user-created virtual router is used in this scenario:  <pre>&lt;Warn:HAL.Ipv6FIB.LongMskUsrVRNoSuprt&gt; MSM-A: Installing IPv6 routes w/ mask &gt; 64 bits with a user virtual router is not supported on G24Xc,G48Xc,G48Tc,10G4Xc,10G8Xc,S-G8Xc, S-10G1Xc and S-10G2Xc. These routes are not installed."</pre>
<b>Summit Series Switches</b>	
xos0056142	In Summit X460-G2 series switches, the software table size is limited by the hardware table size, which limits FDB to 65K.
xos0057210	On Summit X430 series switches, NTP configuration is lost and NTP is not supported when upgrading to ExtremeXOS 15.6.

## Resolved Issues in ExtremeXOS 15.6.4-Patch1-5

The following issues were resolved in ExtremeXOS 15.6.4-Patch1-5. ExtremeXOS 15.6.4-Patch1-5 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.4. For information about those fixes, see the release notes for the specific release.

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

CR Number	Description
<b>General</b>	
xos0061462	Traffic forwarding on specific ports in a VLAN is affected when another port in same VLAN is configured for port specific tag.
xos0062692	Radius-accounting works incorrectly when doing NetLogin MAC and dot1x authentication for the same client.
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.
xos0063457	Configuration for adding network VLAN port in STP for subscriber is not saved.
xos0063478	Traffic drop occurs while adding new member port to the existing LAG group and PSTAG is configured on the port.
xos0063493	Traffic stops after disabling, and then enabling LAG ports in VPLS service VLAN with static FDB.
xos0059481	Static FDB is programmed incorrectly in hardware after a stack failover.
xos0059742	SER not enabled for internal TCAM memories.
xos0063359	The process rtmgr 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.
xos0063547	Process ACL ends unexpectedly after applying a policy file with source zone as a match condition.

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

CR Number	Description
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.
<b>Summit Stack</b>	
xos0061777	Standby nodes do not come back up to operational state after they go into failed state.
xos0063490	CFM stays down after slot reboot on a stack.
<b>Summit X460-G2 switches</b>	
xos0063382	Summit X460-G2 series switches with VIM-2Q modules display error messages "SPI-ROM / MDIO 0:55 Bad Checksum lane 0,1,2,3" during bootup.
<b>BlackDiamond 8800 Series Switches</b>	
xos0063333	In BlackDiamond 8800 series switches, optics information is not detected and ports remain in "Ready" state after reboot.

## Resolved Issues in ExtremeXOS 15.6.4-Patch1-3

The following issues were resolved in ExtremeXOS 15.6.4-Patch1-3. ExtremeXOS 15.6.4-Patch1-3 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.4. For information about those fixes, see the release notes for the specific release.

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

CR Number	Description
<b>General</b>	
xos0057269	SNMP trap extremelpSecurityViolation is sent with incorrect VLAN description.
xos0058555	Inconsistent behavior allowed in secondary EtherType settings on LAG ports when present on both tagged VLAN and tagged VMAN.
xos0059042	Kernel panic occurs during FDB flush operations caused by deleting private VLANs.
xos0061507	SNMPget on EXTREME-SOFTWARE-MONITOR table returns value with incorrect OID.

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

CR Number	Description
xos0061781	Identity manager entries become stale when clients are moved from one port to another in sub-VLANs.
xos0061855	Configured OSPF neighbor is not retained after rebooting.
xos0062240	Port that was administratively disabled becomes up after enabling rx pause.
xos0062366	After rebooting, DHCP binding entries are not restored using vr-default.
xos0062367	ACL process ends unexpectedly on repeated refresh of ACL policy with clear-flow action.
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.
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.
xos0063134	Traffic stops after disabling, and then enabling LAG portst having pstag with static FDB.
xos0063172	ACL action "redirect-port-list" does not take effect when another slice has a rule to match all packets with deny action.
xos0063204	Traffic stops on LAG ports when frequently modifying the sharing group.
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.
xos0063368	In an MLAG configured switch, FDBs are not installed in hardware after reboot if there are frequent MACMoves between MLAG port and ISC.
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".
xos0063506	Traceroute MAC in CFM domain does not return information about destination switch
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.
xos0062255	CEP CVID configurations is missing after adding/deleting the port from sharing.

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

CR Number	Description
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.
xos0063228	Adding new ports to LAGs dynamically (through configure command) drops broadcast/unknown unicast packets between the VPLS peers.
<b>BlackDiamond 8800 Series Switches</b>	
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.
<b>Summit Stack</b>	
xos0063242	Stacks configured as DHCP clients do not respond to ping after failover.
xos0061027	For SummitStacks, creating or deleting non-default QoS profiles may cause some ports to flap.
xos0063344	With MLAG and LAG configurations, when a stack node comes up after a reboot, FDB entries flooded from other slots are programmed on incorrect ports internally.
<b>Summit X460 Series Switches</b>	
xos0063206	Cannot add L2 entries in hardware due to a full L2 table caused by hash collisions.
<b>Summit X670 Series Switches</b>	
xos0063137	Known unicast traffic is not shared between the stacking high-gig trunk ports.



## Resolved Issues in ExtremeXOS 15.6.4

The following issues were resolved in ExtremeXOS 15.6.4. ExtremeXOS 15.6.4 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 12.6.5.2-Patch1-3 , ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.3. For information about those fixes, see the release notes for the specific release.

**Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
xos0050590	Multicast cache entries are not programmed in hardware due to false resource full condition due to IPv6 cache entries.
xos0056885	IGMPv3 packets with invalid checksum are not dropped.
xos0058143	When creating a netlogin-user, some information is truncated in the output of the command.
xos0058775	snmpMaster process ends unexpectedly while fetching system MAC and boot count information simultaneously from switch.
xos0059569	OSPFv3 external routes flap and traffic loss occurs when introducing new ABR which has reachability to destination.
xos0060792	SNMP authentication failure log message and trap is inappropriately generated when switch detects "Not In Time Windows" error.
xos0061505	BGP routes requiring two level of recursive lookup is getting programmed with wrong next hop in hardware after topology change in network.
xos0061788	The process devmgr ends unexpectedly during snmpwalk when continuous EMS logs are sent to the switch console.
xos0062017	DHCP trusted port configuration is getting lost after disable and re-enable of LAG.
xos0062045	LLDP packets are being tunneled over L2VPN
xos0062441	rtMgr process ends unexpectedly when IPv6 static route is deleted
xos0062537	HAL crash occurs when redirect-port-list action contains more than 64 ports.
xos0062914	The process mcmgr ends unexpectedly after receiving corrupted IGMPv3 join packets on MLAG ports.
xos0062494	Source MAC addresses learned through MVRP packets on a blocked port (STP) cause traffic to be dropped.

**Table 7: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
xos0062508	When a port is added in a loopback VLAN, OSPFv3 route is not advertised with /128 mask.
<b>Summit X480 Series Switches</b>	
xos0061071	On Summit X480-48t switches, optic-info is not detected for combo ports.
<b>Summit X670 Series Switches</b>	
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.
xos0062487	QSFP+ optics are detected as unsupported after rebooting.
xos0063052	Traffic loss occurs on computer connected to Summit X670v-48t switches when the connected switch port is oversubscribed in 100 MB mode.
<b>Summit X460-G2 Series Switches</b>	
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.
xos0062867	On Summit X460-G2 series switches, RSVP-TE secondary path is not updated in hardware when the primary goes down.
<b>Summit X440 Series Switches</b>	
xos0062621	On Summit 440-8p switches, the <code>show fan</code> command output displays that the fan is unsupported.
<b>SummitStack</b>	
xos0061067	On SummitStacks with PoE and non-PoE switches, executing the command <code>show inline-power slot &lt;slot_number&gt;</code> produces an XML error.
xos0061791	On SummitStacks containing master and standby nodes of different switches, the standby node may go to failed state after a node reboot.
<b>BlackDiamond 8800 Series Switches</b>	
xos0061748	Error message "aspendiags: Unable to read GBIC EEPROM" appears during bootup of c-series I/O modules.

## Resolved Issues in ExtremeXOS 15.6.3-Patch1-9

The following issues were resolved in ExtremeXOS 15.6.3-patch1-9. ExtremeXOS 15.6.3-patch1-9 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.3. For information about those fixes, see the release notes for the specific release.

**Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
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.
xos0057584	When there are equal cost ASBR routes, OSPFv3 is sharing only one path to route manager.
xos0059165	OSPFv3 crash occurs after deleting newly added VLANs.
xos0059560	After reboot, OSPFv3 fails to select the best path to destination.
xos0062719	Allow use of 3rd-party optics without any addition license.
xos0062728	OSPFv3 best path is not selected after issuing <code>restart ports a11</code> in peer switch.
xos0062789	Disabling learning on LAG ports does not flush FDB entries.
xos0057538	OSPFv3 fails to select the best cost external route.
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 a11</code> is executed in area border router.
xos0062710	On BlackDiamond 8800 or BlackDiamond X8 series switches, with distributed IP ARP mode on, ECMP routes are sometimes not installed when gateways flap.
xos0062824	Extra ACL slice allocated for QoS even with default configuration.

**Table 8: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>Summit X430 Series Switches</b>	
xos0059486	On Summit X430 series switches, optics are not detected after repeated removal and reinsertion of optics when CPU is busy.
<b>Summit X440 Series Switches</b>	
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.
<b>SummitStack</b>	
xos0062640	snmpMaster process ends unexpectedly when SNMP set request is received during master-backup sync operations.
xos0062800	Stack node fails because of license mismatch for 3rd-party optics.
<b>BlackDiamond 8800 Series Switches</b>	
xos0055892	Errors appear indicating that the SFP/GBIC EEPROM cannot be read for some SFPs optics:  <Erro:Card.Error> Slot-5: aspendiags: Got bad rv (-1) reading XFP Diagnostic Monitor EEPROM for slot 5, port 8  <Erro:Card.Error> Slot-5: aspendiags: Unable to open XFP EEPROM for port 8, rv -1
xos0062805	With distributed IP ARP mode on, table full warning messages appear on BDX XL and BD8900 XL I/O modules when the number of ARPs attached to one XL slot exceeds 50% capacity (greater than 32,000 on BDX XL or greater than 8,100 on BD8900 XL).
<b>BlackDiamond X8 Series Switches</b>	
xos0054860	Configuring and unconfiguring large ACL policy files two or three times causes the ACL to fail during subsequent installs.
xos0056773	Resetting a connected I/O fabric module, produces the following error messages:  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)
xos0060426	"Unable to read QSFP info" error messages appear when reading EEPROM from optics.

## Resolved Issues in ExtremeXOS 15.6.3-Patch1-8

The following issues were resolved in ExtremeXOS 15.6.3-patch1-8. ExtremeXOS 15.6.3-patch1-8 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.3. For information about those fixes, see the release notes for the specific release.

**Table 9: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
xos0058135	Tagged VLAN traffic is inappropriately grabbed by VMANs after removing service VLANs from VPLS instances. This happens when the same port is tagged in the service VLAN and untagged in the VMAN.
xos0061265	The <code>unconfigure switch erase all</code> command does not erase information in NVRAM for Summit G2 series switches.
xos0062113	The <code>show power</code> command output does not display power usage for PSUs with part numbers starting with "800515".
xos0062705	Kernel oops can occur after clearing IPMC FDB in a stack.
xos0053828	HAL process ends unexpectedly when all entries from network-zone are deleted and associated ACL is refreshed.
xos0058514	SnmpEngineBoots value appears as a negative value on certain switches.
xos0062427	EDP process ends unexpectedly when CDP packets without portId TLV are received.
xos0058808	Rarely, MAC addresses of authenticated clients learned on NetLogin-enabled ports are not programmed in hardware.
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.
xos0062271	CLI memory leak occurs when executing show commands with include option through script.
xos0062380	Switch rejects incorrect LSP configurations as expected, but this operation still uses LSP indexes in hardware.
xos0062472	Source MAC addresses learned through CDP packets received on EAPS-blocked ports cause traffic to be dropped.
xos0061886	SNMP master process ends unexpectedly with signal 6 with certain sequence of <code>snmpbulkget</code> and <code>snmpget</code> .

**Table 9: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>BlackDiamond X8Series Switches</b>	
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 "Detected Tx unit hung" error messages.
xos0062499	Multicast packets are dropped in Layer 2 bridged VLANs.
<b>SummitStack</b>	
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.
xos0062522	In Summit stacking switches, standby slots go to failed state when very large log messages are continuously generated in the switch.
xos0062570	In SummitSacks, executing the command <code>enable sflow ports all</code> enables sFlow inappropriately on stacking ports.
<b>Summit X460-G2 Series Switches</b>	
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.
xos0059763	In Summit X460G2-24x-10G switches, ports on VIM-2x sometime remain in ready state after rebooting.
xos0061486	Combo ports have unsupported autonegotiation and half-duplex settings.
<b>Summit X670 Series Switches</b>	
xos0061167	Links become active without a connection with tri-speed Base-T SFP installed.

## Resolved Issues in ExtremeXOS 15.6.3-Patch1-5

The following issues were resolved in ExtremeXOS 15.6.3-patch1-5. ExtremeXOS 15.6.3-patch1-5 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.3. For information about those fixes, see the release notes for the specific release.

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

CR Number	Description
<b>General</b>	
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.
xos0057464	TLS protocol is impacted by CRIME (Compression Ratio Info-leak Made Easy) vulnerability.
xos0058842	Tagged VLAN traffic is dropped at ingress of 100G4X-XL module/Summit X460-G2 port when the same port is added to another CEP-VMAN VPLS.  This issue does not occur on non-100G4X-XL modules/ports on BlackDiamond X8 series switches.
xos0059942	SSH connection ends when show commands produce lengthy output.
xos0061092	Traffic forwarding on VPLS-serviced VMAN stops after link flap.
xos0061198	Disabling VPN-VRF affects traffic on another VPN-VRF.
xos0061379	Switch temperature value retrieved using SNMP get operation is incorrect.
xos0061385	EAPS process ends unexpectedly after deleting EAPS shared-port configuration.
xos0061822, xos0061957	HAL process ends unexpectedly during failover when switches have ACL policies without meter action.
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.

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

CR Number	Description
<b>BlackDiamond 8800 Series Switches</b>	
xos0061994	Packets are not forwarded over VPWS after rebooting BlackDiamond 8800 series switches.
xos0062060	FDB entries are not programmed in hardware even though hardware resources have sufficient capacity.
<b>BlackDiamond X8 Switches</b>	
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.
<b>Summit Family Switches</b>	
xos0058790	SummitStacks are rebooted after executing <code>debug hal show optic-info slot &lt;master slot#&gt; port &lt;port#&gt;</code> from backup node.
<b>Summit X460 Series Switches</b>	
xos0059688	Default debounce timer value (in ms) for alternate stacking port should be "0" instead of "150" in Summit X460-48p switches.
<b>Summit X670-G2 Series Switches</b>	
xos0061661	Multicast add failures occur on Summit X670-G2 stacks with the following error message:  <Error:Kern.IPv4Mc.Error> Slot-1: Unable to Add IPmc sender entry s,G,v=192.1.2.63,224.0.0.5,291 IPMC 974 flags 4 unit 0, Entry exists



## Resolved Issues in ExtremeXOS 15.6.3-Patch1-3

The following issues were resolved in ExtremeXOS 15.6.3-patch1-3. ExtremeXOS 15.6.3-patch1-3 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.5.2-patch1-3, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.4.2, and ExtremeXOS 15.6.3. For information about those fixes, see the release notes for the specific release.

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

CR Number	Description
<b>General</b>	
xos0053821	IPv6 neighbor advertisements for VRRP virtual IP address uses virtual MAC address as source MAC address instead of switch MAC address.
xos0055376	After executing the commands <code>clear netlogin state</code> and <code>clear fdb</code> , NetLogin MAC address authenticated entries are not blackholed.
xos0055514	Process "bcmRX" consumes 25% of CPU with one-to-many mirroring feature enabled.
xos0055945	Error message "exosmc: ip_mc_handle_msdp_data:2038: MC: Ingress vif not found" appears when sending multicast-tagged packets to member VLAN of translation-VLAN/ subscriber VLAN of private VLAN.
xos0056243	OSPF process ends unexpectedly during frequent route re-calculation caused by switch reboot or MSM failover or BFD flap events.
xos0056340	Unknown Layer 2 traffic from Isolated subscriber VLANs are forwarded to the remote MLAG ports, even though local MLAG ports are up.
xos0056553	The output of the command <code>show fdb netlogin all</code> does not show 's' or 'd' flags for NetLogin entries.
xos0058188	LACP member ports in the remote end are removed from the aggregator when back-to-back MSM failovers are executed in the local end.
xos0058671	In VPLS, traffic destined to unknown MAC addresses is duplicated after changing dot1q tag include or exclude configuration.
xos0059266	VPLS traffic is not switching over to alternate RSVP LSP after disabling the active LSP.
xos0059655	Error "Unable to delete permanent entry" appears while deleting static blackhole FDB entries.

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

CR Number	Description
xos0059730	The process mcmgr ends unexpectedly after removing a slot from existing SummitStack and unconfiguring that slot using the command <code>unconfigure slot slot_number</code> .
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.
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.
xos0060354	ExtremeXOS ScreenPlay using IPv6 does not work with HTTPs.
xos0060449	NetLogin MAC-based mode does not work as expected with PVLANS.
xos0060621	Process rtmgr ends unexpectedly after unconfiguring, and then re-configuring IP addresses. Issue happens only when switches have saved configurations where the IP route has the default gateway as local interface address, which is invalid.
xos0060716	Need support for new ACL action "redirect-vlan" to redirect matched packets to all ports in specified VLANs.
xos0060780	With VRRP enabled, local VLAN's direct route is not installed in hardware after reconfiguring the VLAN's IP address.
xos0060794	VRRP advertisement interval configuration changes after upgrading ExtremeXOS from 12.6 to 15.4 or later releases. Issue occurs only when interval is configured in milliseconds.
xos0060909	In UPM profiles the variable EVENT.TIME incorrectly has the current time rather than the time when the event was queued/triggered.
xos0061009	The output of the command <code>show netlogin MAC</code> output displays username for unauthenticated client.
xos0061038	Loops occur in EAPS-protected VLANs, after peer reboot, if a VLAN's port is also protected by ELSM.
xos0061069	In Netlogin ISP mode, client MAC addresses configured as static FDBs are removed after reboot.
xos0061085	Kernel oops occurs while deleting VR with enable BGP export and IPARP proxy configurations.
xos0061178	Dynamic ACL for gratuitous ARP violation on LAG member ports are incorrectly getting installed on LAG master ports.
xos0061222	Gratuitous ARP packets for VRRP virtual IP addresses have ARP sender addresses as physical MAC addresses, instead of VRRP virtual MAC addresses.
xos0061565	The TCL function, "clock scan," generates errors with default time zone configuration.
xos0061586	In two-tier MLAG topology, loops occur if the VLAN header has the CFI bit set.

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

CR Number	Description
xos0061699	Traffic is dropped when moving idmgr client from one port to another with role-based authentication.
xos0061835	The process exsh causes excessive CPU utilization after performing continuous Telnet/SSH for the switch.
xos0061922	Dynamic ACLs applied as "any" fail to install in hardware after upgrading ExtremeXOS from any release other than EXOS 15.3.
xos0054039	IP multicast traffic is not forwarded on PSTAG VLANs when it shares ports with other IGMP snooping-enabled VLANs or other L3 VLANs.
xos0055398	Authentication fails for a Netlogin client in dot1x mode, since the port added untagged in one VLAN cannot be moved to another VLAN.
xos0056263	The following error message occurs when deleting VPLS instances: <pre>"&lt;Error:Kern.MPLS.Error&gt; : bcm_custom_extr_vfp_tagged_vlan_port_add, Entry exists, rv 0"</pre>
xos0057407	Hops fields in DHCP packets are not incremented when processed by Bootprelay.
xos0060407	After disabling, and then enabling EBGp, new routes from different autonomous systems (ASs) are not considered for best path calculation.
xos0061517	LACP adjacency fails while forwarding the PDU with l2pt profile over L2VPN tunnels when MPLS PHP is enabled.
xos0061656	Nodes remain in the "FDBSync" state due to temp-flooding while rebooting the stack.
xos0061965	Configuring ESRP member VLANs (VRRP-enabled) produces errors.
<b>BlackDiamond 8800 Series Switches</b>	
xos0060891	While sending continuous join and leave as well as multicast streams multiple times allowing the streams to get aged out, the kernel error messages are seen for add/delete operation with reason "Entry not found"
xos0061338	Packets are not switched on port-specific, tag-enabled ports on XL series I/O modules.
xos0061796	Error message "aspenSmIpmcAddEgressPort: group does not exist" appears during switch reboot or MSM failover.
xos0057624	Traffic loss occurs on PVLAN after restarting VRRP process in VRRP master switch.

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

CR Number	Description
<b>BlackDiamond X8 Switches</b>	
xos0057827	Layer 2 multicast traffic is not forwarded to IGMP receivers after MM (management module) failover.
xos0059156	VRRP control packets are dropped due to congestion in tx queue under scaled environments.
xos0059603	In BlackDiamond X8, Management Modules fail to complete booting when there is a failed Fabric Module in the chassis.
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%.
xos0061902	BlackDiamond X8 series switches use VLAN instance as index instead of router interface (rtif) for ARP entries.
xos0058661	ARP reply packets are dropped with continuous IP unicast miss packets with DSCP 48.
<b>Summit X430 Series Switches</b>	
xos0061864	FDB process consumes more than 20% utilization when Summit X430 switches are configured with 100+ VLANs.
<b>Summit Family Switches</b>	
xos0060965	Netlogin process ends unexpectedly when web-based users log out and then refresh the logout window a few times.
xos0061191	In SummitStacks with PVLAN configuration, after stack failover, VRRP advertisement has incorrect VLAN ID.
xos0057438	Memory depletion occurs in Backup/Standby nodes of SummitStack with highly scaled IPFIX flow records.
xos0058819	Part information for power supplies of stack nodes is not visible from master nodes.
<b>Summit X440 Series Switches</b>	
xos0059932	ACL slice utilization is greater with default configuration causing ACL installation failure with user ACLs due to limited resources.
<b>Summit X460 Series Switches</b>	
xos0060517	Untagged ports in service VMANs take other VLANs' traffic after events such as add/delete port in VLANs/VMANs, change tag in VLANs/VMANs, etc.
xos0061180	In Summit X460 stack with mixed alternate and native stacking enabled slots, traffic ingressing one specific slot is not forwarded to other slots.

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

CR Number	Description
<b>Summit X460-G2 Series Switches</b>	
xos0060559	Hot-swapping fan modules in Summit X460-G2 series switches puts some fans in failed state.
xos0060998	MAC-based authenticated sessions over RADIUS are erroneously deleted by FDB.
xos0061313	Summit X450/X460-G2 series switches 10G ports flood unicast traffic to other ports on the switch.
<b>Summit X670 Series Switches</b>	
xos0061730	Traffic is affected when unconfiguring ACL on another port.
xos0061770	Detected parity errors may cause a kernel crash.
<b>Summit X670-G2 Series Switches</b>	
xos0060948	Kernel crashes on Summit X670G2-48x-4q switches when stack ports are configured for v80 stacking.
xos0061818	After rebooting Summit X670-G2 stacks, links come up after long delay of ~20 minutes when QoS profile is configured on 500+ VLANs.

## Resolved Issues in ExtremeXOS 15.6.3

The following issues were resolved in ExtremeXOS 15.6.3. ExtremeXOS 15.6.3 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.4.6-Patch1-13, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.3.4-patch1-5, and ExtremeXOS 15.6.2. For information about those fixes, see the release notes for the specific release.

**Table 12: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
xos0060825	Double-tagged CFM frames are dropped by kernel in VMAN environment.
xos0060736	gPTP propagation delay is not calculated correctly and ports become AVB incapable.

## Resolved Issues in ExtremeXOS 15.6.2-Patch1-6

The following issues were resolved in ExtremeXOS 15.6.2-patch1-6. ExtremeXOS 15.6.2-patch1-6 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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.4.6-Patch1-13, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.3.4-patch1-5 and ExtremeXOS 15.6.2. For information about those fixes, see the release notes for the specific release.

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

CR Number	Description
<b>General</b>	
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.
xos0060711	TRILL port metrics are incorrect when TRILL is initially disabled during startup and then enabled later.
xos0054348	Cannot delete flow names after deleting, and then creating, the flow while the ACL is installed.
xos0054605	Bi-Directional L2 traffic fails between TRILL routing bridges after enable/disable sharing on a network vlan link.
xos0058880	Packets are not switched to primary path after recovering from path failure in MPLS RSVP-TE.
xos0059574	OSPF packets larger than 8,192 are dropped even with jumbo frame enabled.
xos0060103	SNMP walk does not return all VLANs under extremePortVlanStatsTable.
xos0060228	HAL process ends unexpectedly in rare circumstances while rebooting switch with default speed configuration on 10G ports.
xos0060613	Null adjacency error message appears after rebooting Leaf node in TRILL.
xos0060756	Spine node dropping TRILL packets received from Leaf node after a link flap.
xos0059288	Cannot install multiple XMOD files in BlackDiamond 8800 and BlackDiamond X8 switches due to lack of space.
xos0058221	Rarely, OSPFv3 process ends unexpectedly with signal 11 when link flaps occur.
xos0060968	TRILL process ends unexpectedly on Leaf node after disabling, and then enabling TRILL on Spine node.,
xos0060817	Kernel error messages appear on TRILL Spine node after rebooting peer Spine node.

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

CR Number	Description
xos0060703	HAL process ends unexpectedly on TRILL Leaf node after rebooting Spine node.
xos0060615	With TRILL configured, kernel error messages appear on Leaf node after rebooting the Spine node.
xos0060612	With TRILL configured, kernel error messages appear on Spine node after rebooting the Leaf nodes.
xos0049802	Enabling SNTP should be disallowed on switches with NTP already enabled.
xos0051961	Unable to block IPv6 traffic from SSH/Telnet/Web interface by access-profile policy.
xos0053634	MAC-lockdown-timeout on user ports does not work as expected, if Netlogin is enabled on those ports.
xos0057211	Traffic gets forwarded for blackholed MAC address when limit learning is enabled.
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.
xos0060088	Kernel oops triggered rarely during continuous addition/deletion of ARP entries for long duration in presence of high CPU utilization.
xos0060119	Changing the primary TACACS server configuration locks out TACACS-authenticated users.
xos0060615	With TRILL configuration, kernel error message seen on Leaf node after rebooting the Spine node.
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.
xos0058994	PoE is not delivering power to several model phones when legacy mode is enabled.
xos0059989	Configuring non-persistent command using UPM script shows dirty bit(*) in the prompt.
xos0060092	Fetching values using SNMP for "extremePortQosStatsTable" does not work correctly.
xos0060176	The process <code>rtmgr</code> ends unexpectedly with signal 11 because of segmentation fault. This occurs only when default route is exported from BGP neighbor.
xos0060214	Process <code>netTools</code> ends unexpectedly during reboot of switch with VRRP track-ping configuration.
xos0060417	The command <code>show trill neighbor</code> does not show all the neighbor information.

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

CR Number	Description
<b>Summit Family Switches</b>	
xos0059950	In Summit series switches, you cannot download bootROM images from memory card.
xos0059462	Timezone configuration is not applied to standby nodes after stack reboot.
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.
<b>Summit X430 Series Switches</b>	
xos0059934, xos0057028	For Summit X430 series switches, downloading BootROM corrupts the bootloader.
<b>Summit X770 Series Switches</b>	
xos0055746	Stacking port link flap occurs on Summit X770 series switches when using 3-meter QSFP+ cables.
<b>Summit X460 Series Switches</b>	
xos0059671	On Summit X460 series switches with 750 W power supplies installed, log messages "Power usage data unknown" appear.
xos0059320	CCM is dropped for "Hardware Down MEPs" when they are received on ports that are blocked by ERPS.
<b>BlackDiamond X8 Series Switches</b>	
xos0060210	Rarely, HAL process may end unexpectedly due to buffer overflow condition while running diagnostics for Management Module.
<b>BlackDiamond 8800 Series Switches</b>	
xos0060301	Rarely, ports go into ready state when the connected devices are continuously auto-negotiating to different speeds. Disabling/enabling such port can trigger I/O module reboots.



## Resolved Issues in ExtremeXOS 15.6.2-Patch1-1

The following issues were resolved in ExtremeXOS 15.6.2-patch1-1. ExtremeXOS 15.6.2-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 12.6.5.2-Patch1-3, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5.4-Patch1-9, ExtremeXOS 15.2.4.5-Patch1-5, ExtremeXOS 15.3.1.4-patch1-47, ExtremeXOS 15.3.4.6-Patch1-13, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.3.4-patch1-5 and ExtremeXOS 15.6.2. For information about those fixes, see the release notes for the specific release.

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

CR Number	Description
<b>General</b>	
xos0059077	Getting error after executing the <code>upload log</code> command multiple times.
xos0059789	Dos-Protect ACL is not cleared after continuous DOS attack occurs.
xos0059945	Memory corruption occurs rarely due to packet buffer overrun while sending/receiving control packets between slots.
xos0056913	OSPFv3 process ends unexpectedly when link goes down on set of ports in the switch.
xos0058391	Switches allow installing ACL policy with meter even though the corresponding meter is not yet created.
xos0059581	Rtmgr process ends unexpectedly when OSPF external routes are deleted from the route table.
xos0059584	OSPFv3 process does not re-advertise IPv6 routes to rtmgr process even after receiving Link LSA from peer. This happens when the first advertisement was sent before receiving Link LSA from peer.
xos0060100	Kernel oops occurs due to memory corruption caused by slow-path forwarded traffic.
<b>Summit Family Switches</b>	
xos0056230	SNMP query on "extremeMemoryMonitorsystemTable" does not show backup information, if slot2 is master and slot1 is backup.
xos0056342	Misleading power supply unit (PSU) traps are sent when PSUs are inserted or powered on/off.
xos0058460	Unable to configure multiple LSP from core node to spoke node in Summit X460-G2 and X670-G2 series switches.
<b>E4G-200 Cell Site Routers</b>	
xos0058239	In E4G-200 cell site routers, power supply status displays incorrect value in the output of the <code>show power</code> command.

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

CR Number	Description
<b>Summit X440 Series Switches</b>	
xos0058300	Packets are dropped on combo ports when the preferred medium is configured as copper force.
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.
<b>Summit X460-G2 Series Switches</b>	
xos0055648	I2C error messages appear during boot up of Summit X460-G2 stacks.
xos0059827	During Summit stack failover, Kernel oops appears infrequently, caused by corruption in VR ID of resolved ARPs.
<b>BlackDiamond 8800 Series Switches</b>	
xos0059648	Static ARP entries are not properly synced with new Master Switch Fabric Module after failover.
xos0059605	Sys-health-check output shows false fabric port flap events between Master Switch Fabric Module (MSM) and I/O module.

## Resolved Issues in ExtremeXOS 15.6.2

The following issues were resolved in ExtremeXOS 15.6.2. ExtremeXOS 15.6.2 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, and ExtremeXOS 15.6.1. For information about those fixes, see the release notes for the specific release.

**Table 15: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
xos0054199	Ingress traffic stalls on port when switches receive continuous 802.3x pause frames on egress ports for that traffic stream.
xos0058833	L2PT over VPLS is not working when protocol filter is added to profile with a CoS value.
xos0056228	TFTP get operation fails when the remote file exists in second level sub-directory.
xos0058120	After running the command <code>debug cli run python exosjson.py getnext mod.struct</code> , running the command <code>show configuration</code> produces the error: "couldn't open './config/xos_config.xml': no such file or directory."
xos0058203	The command <code>debug cli run python</code> allows a user-created Python script (stored in <code>/config</code> , <code>/usr/local/cfg</code> ) to run with root rights posing a security risk.
xos0058393	The TCL command <code>clock format</code> is not available for CLI scripting and attempting to use it produces an error.
xos0058695	Process <code>emsServer</code> ends unexpectedly with signal 6 when multiple VRRP messages are logged.
xos0058968	Error log "Function Pointer Database is not fully initialized" appears during bootup on non-Summit platforms.
xos0059037	Pre-emphasis <code>show</code> command displays incorrect values for non-Summit X460 series switches' slots in mixed stacks.
xos0059243	The process <code>exsh</code> ends unexpectedly after executing a <code>show</code> command with a port list followed by invalid letters (for example, <code>show port 1:1,1:2ab</code> ), and then pressing <b>TAB</b> .
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.
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.

**Table 15: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>BlackDiamond X8 Series Switches</b>	
xos0055433	The tDiag process occasionally ends unexpectedly after executing <code>show debug system-dump MM B</code> from MM-A, when MM-B does not contain system dump.
xos0057352	Kernel crash occurs when there is a Layer 2 loop in the network.
xos0057560	Accessing ScreenPlay while running a script can cause the thttpd process to end unexpectedly with the following error: <code>&lt;Error:DM.Error&gt; MM-A: Process thttpd Failed MM-A rebooted</code>
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 12,000 IPv4 routes and above 32 IPv4 local hosts.
xos0058568	Some front panel ports cannot be enabled after rebooting the I/O module.
xos0059104	ACL policies are not installed in hardware after management module failover.
xos0059343	The process snmpMaster might end unexpectedly during upgrade from ExtremeXOS 15.3 to 15.5 for some SNMP community names.
<b>Summit Series Switches</b>	
xos0056971	In Summit 670G2-72x and Summit 460-G2 stacks, executing a failover produces the following Kernel error: <code>&lt;Error:Kern.Card.Error&gt; Slot-1: KICM: Could not change SP1 type for NONE neighbor. idx1=0,idx2=0</code>
xos0058537	Switches become unresponsive and drop traffic when they have a high number of traffic streams and AVB enabled ports.
xos0059447	Can use Python scripts to access debug shell and execute commands even though debug mode is not enabled making switches vulnerable to unauthorized use.
<b>SummitStack</b>	
xos0055287	On Summit X480-24x(SSV80)/X670v stacks, MPLS process ends unexpectedly with signal 6 when enabling stacking using the command <code>configure stacking easy-setup</code> in the switch with core and MPLS feature pack license. This issue does not occur with Summit X460 and X440 stacks.
xos0057767	Static FDB associated with VPLS service VLAN is not programmed in hardware after reboot when "disable learning" is configured.
<b>Summit X430 Series Switches</b>	
xos0059524	Link status is incorrect when auto-polarity setting is off.

Table 15: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)

CR Number	Description
<b>Summit X440 Series Switches</b>	
xos0058068	Summit X440-24tDC switches are reporting maximum temperature limit 60°C under normal conditions.
xos0058301	In Summit X440 series switches, error message "mounting /dev/hda4 on /data failed" appears during bootup.
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").
<b>Summit X450 Series Switches</b>	
xos0057647	Packets are forwarded to CPU after deleting the VLAN with <code>disable learning</code> .
<b>Summit X460 Series Switches</b>	
xos0058589	SummitStack reboots due to temperature out of range messages.
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.
<b>Summit X460-G2 Series Switches</b>	
xos0058896	Running slow path traffic for long durations causes Summit X460G2 stack instability.
xos0055189	In Summit X460-G2 stacks, the command <code>show power</code> fails to display power usage and produces the error "Failed reading Slot-B power on time" during slot reboot.
xos0058209	In Summit 670G2-72x and Summit 460-G2 stacks, executing a failover produces the following Kernel error:  <Error:Kern.Card.Error> Slot-1: KICM: Could not change SP1 type for NONE neighbor. idx1=0,idx2=0
xos0058366	Error message "phy_state_machine: PHY_NOLINK link=0 autoneg=1 fallback=0" appears when management port is not active.
xos0059577	On Summit X460G2 series switches, can't install ExtremeXOS SSH XMOD image.
<b>Summit X670 Series Switches</b>	
xos0059128	In Summit X670 series switch, all LEDs are blinking at a faster rate.
<b>Summit X670-G2 Series Switches</b>	
xos0059445	Link flaps occur when stacks are firmed with 3 m/5 m QSFP+ passive copper cables.
xos0058674	When 10G (SFP+) end of 5 m passive copper fanout cable is connected to 10G ports of Summit X670-G2 series switches, link flap occurs.

**Table 15: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>Summit X770 Series Switches</b>	
xos0059573	Factory installed image incorrectly references "x450" in the image name.
<b>ACL</b>	
xos0056423	The command <code>show access-list meter port</code> does not display the meters applied on the port via policy.
xos0057328	ACL rule to match IPv6 packets with arbitrary mask is not working as expected.
xos0058810	Both <code>show access-list usage acl-slice</code> and <code>debug hal show device acl-slice slot &lt;slot&gt; unit &lt;unit&gt;</code> commands show the same output.
xos0059330	With dual master switch fabric module (MSM) installed, clear-flow ACL intermittently fails.
<b>AVB</b>	
xos0057964	ACL process ends unexpectedly when accessing AVB related dynamic rule from ExtremeXOS ScreenPlay.
xos0058603	PTP follow up does not happen correctly when correction field is greater than 32 bits.
<b>FDB</b>	
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.
<b>DHCP</b>	
xos0055108	The bound IP address is not being reflected in the command <code>show vlan</code> .
xos0057976	DHCP bootprelay option 82 with aggregation port shows incorrect circuit ID.
xos0058585	DHCPv6 relay is not working with Smart Relay Mode On (both parallel and sequential). DHCPv6 requests are not relay forwarded causing clients to not get IPv6 addresses.
<b>ERPS</b>	
xos0058464	In ERPS rings, blocking the control channel by deleting the ports from the control VLAN causes a short loop in the ring.
xos0056122	Configuring MEP IDs through ERPS (for dynamic CFM creation) commands are unavailable.  However, existing configuration files that contains ERPS commands with east/west MEP IDs specified are still loaded correctly, with the underlying CFM created properly by ERPS. Also, the command <code>show configuration</code> displays the configuration properly. However, you cannot modify the MEP ID through the CLI.
xos0057179	ESRP feature is not enabled immediately after the installing an Advance Edge license.

**Table 15: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>IP Routing Protocols</b>	
xos0052696	Static routes are installed in route manager when the next hop is configured as its own loopback or as a local VLAN interface IP address.
xos0057672	The process rtmgr ends unexpectedly when disabling GRE tunnels.
xos0058056	OSPF-opaque, LSA-related configurations do not appear in output of the <code>show configuration</code> command.
xos0058611	OSPFv2 external routes are not updated in routing tables after disabling uplink ports in peer switches. During this condition, the route to Advertising Router (external routes) is available by OSPF neighbors, but external routes are not updated dynamically in the routing tables.
xos0058683	RIP packets are dropped when another VLAN has a secondary IP address configured.
xos0058801	IPv4 ECMP route entries learned by a routing protocol are sometimes removed from hardware when one of the next hop gateways goes down, but other gateways remain up.
xos0059305	OSPF consumes a large amount of memory when a large number of Link State Acknowledgment packets are queued up for transmission.
<b>MLAG</b>	
xos0058873	FDB entries are learned incorrectly on VMANs in MLAG peers when the MLAG ports are CEP ports for multiple VMANs with different CVIDs.
<b>sFlow</b>	
xos0053584	sFlow displays incorrect VLAN tag when collecting on tagged VMAN ports.
xos0059222	SFLOW-sampled packets are flooded out of VLANs when these same packets are software learned.
<b>Security</b>	
xos0057601	The command <code>disable ip-security arp learning learn-from-arp vlan &lt;vlan_name&gt; port &lt;portname&gt;</code> is not saved/applied after reboot.
<b>STP</b>	
xos0059002	Checkpoint errors occur during execution of STP debug command if switch contains many STP-enabled VLANs.
xos0057785	STP domain tag is removed when all ports are deleted from STP auto-bind enabled-VLANs.
<b>VLANs</b>	
xos0057435	Packets are dropped when learning is disabled in a VLAN when its associated ports are configured with limit learning in another VLAN.

**Table 15: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>VPLS</b>	
xos0056994	Unable to add EAPS shared ports to VLANs even after disassociating them from VPLS domains.
xos0058397	Unknown unicast traffic via VPLS is duplicated after back-to-back LSP failover.
xos0058717	The message "Warn:MPLS.RSVPTTE.InternalProb" appears after disabling ports in RSVP secondary path.

## Resolved Issues in ExtremeXOS 15.6.1

The following issues were resolved in ExtremeXOS 15.6.1. ExtremeXOS 15.6.1 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, and ExtremeXOS 15.5.2. For information about those fixes, see the release notes for the specific release.

**Table 16: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>General</b>	
xos0052365	Unconfiguring L3VPN and then configuring it again results in BGP failing to advertise the routes.
xos0053543	Switch responds to a GARP attack for an IP address that is configured as a static IP ARP entry with the switch MAC address, but the switch should respond with the MAC address listed in the static IPARP entry.
xos0054120	Enabling IPv6 smart relay at VR level does not forward the relay-forward solicit from having multiple BOOTPv6 relay enabled agents.
xos0055055	CliMaster process ends unexpectedly when the telnet session disconnects from the switch while it is printing a lengthy debug output.
xos0055311	IPARP entries are cleared immediately after they are learned from source VLAN port when IPARP timeout is disabled.
xos0055476	DHCP decline packets are dropped if the client address field within the DHCP decline packet is 0.0.0.0.
xos0055477	Switches don't allow configuring L2PT protocol filter with dest-mac, etype, etc.
xos0055526	The command <code>set var emp3 \$tcl)</code> produces the error "Error: Unsupported TCL Command", but it should show "% Missing close parenthesis at '^' marker."



**Table 16: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
xos0055585	Multicast traffic can take up to 60 seconds to recover when an ingress port on a first-hop router (FHR) is disabled.
xos0055611	PIM does not failover to alternate source received from the MSDP peers when the primary source fails.
xos0055662	Delays in triggering a cyclic UPM profile, cause the next FireTime to be calculated using the delayed current time.
xos0055685	UPM process ends unexpectedly when the command <code>show config upm</code> is executed repeatedly for a prolonged time.
xos0055754	Login authentication event is generated for Lawful Intercept user if the user logs on using SSH.
xos0055783	Switches drop packets and display "Invalid MAC Binding" error when you remove an existing client and connect a different client with the same IP address.
xos0055849	By enabling ACL log filters, admin user can see the dynamic ACL binding/unbinding information of Lawful Intercept user.
xos0055884	Unicast traffic is dropped completely in fabric after destination MAC addresses are moved from one slot to another.
xos0055939	EXSH assert failure occurs when executing commands containing port lists separated by commas when last port number is followed by TAB key sequence.
xos0056115	Some VRRP VLANs undergo state change when ports belonging to VRRP VLANs that are VLAN/route tracked are disabled/enabled.
xos0057088	Cannot log on using SSH with a 32-character or greater password. After eight logon attempts, no more SSH connections are permitted.
<b>BlackDiamond 8800 Series Switches</b>	
xos0052294	A port in the ISC VLAN is incorrectly allowed to be configured as an MLAG port.
xos0052349	Process rtmgr ends unexpectedly with signal 11 while disabling an I/O slot in a BGP configuration.
<b>BlackDiamond X8 Series Switches</b>	
xos0050424	Slot goes into failed state after running extended diagnostics.
xos0053683	DHCP packets from sub-VLANs are not egressing through the super VLAN ports.
xos0053879	The error "<Error:HAL.MPLS.Error> MM-A: ILM instance 872677376 not found" appears after issuing the command <code>clear counters mpls rsvp-te</code> .
xos0055282	Strict-Priority (SP) and Weighted-Deficit-Round-Robin (WDRR) scheduling are not supported on the BlackDiamond XB-100G4X in this release.  Weighted-Round-Robin (WRR) scheduling is supported on the BlackDiamond X8-100G4X switches in this release.
xos0055644	Add Extreme Networks' serial number to the "easily readable" section of CFP2 optic command <code>debug hal show optic-info slot &lt;slot&gt; port &lt;tpport&gt;</code> .

**Table 16: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
xos0055715	VPLS terminated traffic on 100G slots is not distributed onto LAG ports on BDXA-10G48X modules.
xos0055730	On BlackDiamond X8 switches with BDX8-100G4X modules, MAC addresses are no longer learned on EAPS ring ports after ports are disabled and enabled multiple times.
xos0055733	VPLS: Multicast traffic is dropped on BDX8-100G4X modules when the service port is part of a VLAN with IGMP snooping enabled.
xos0055803	Backup management cards do not enter "In Sync" state if port-specific tags are configured.
<b>Summit Series Switches</b>	
xos0049797	ERPS process ends unexpectedly with signal 11 after issuing the command <code>delete erps submetro</code> .
xos0053227	VLAN statistics are not working after a VLAN with base PSTAG is removed and added back to ports.
xos0054086	VLAN aggregation configurations are not removed after rebooting the switch when configuring a dynamic VLAN as the sub-VLAN.
xos0055623	MLAG ports configured with LACP go down after performing a stack failover.
<b>SummitStack</b>	
xos0053970	NTP process ends unexpectedly during switch reboot.
xos0055980	In Summit-stack, traffic passing through a port from backup node stops after adding that port to another VLAN with port specific tag configuration.
<b>Summit X460 Series Switches</b>	
xos0058043	MSRP: packets are dropped when bandwidth is increased to 611 M.
xos0052683	Stacking does not come up when trying to use one port in alternate and another port in native mode for stacking.
xos0055416	On Summit X460 series switches, after enabling jumbo frames on 10G ports on XGM modules, doing either a save and reboot or changing port speed settings removes the jumbo frames setting.
<b>Summit X480 Series Switches</b>	
xos0057602	Ping across VPLS domain does not succeed if the ingress port in CPE switch is part of both service VLAN (tagged) and service VMAN (untagged).
xos0056976	On Summit X480 series switches with 40G4X VIM modules, FDB entries are not in sync across units when FDB entries are learned and aged out frequently.
xos0057454	FDB entries are not aged out after two continuous MAC address moves.
xos0057752	On Summit X480 series switches when untagged VLAN packets are received on tagged VMAN ports, which in turn receive from untagged VMAN ports bursts of more than 500 packets in line rate, some FDB entries are not learned in hardware.

Table 16: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)

CR Number	Description
<b>Summit X670 Series Switches</b>	
xos0048730	After MLAG is established, enabling, and then disabling, MLAG ports produces error messages: <pre>"07/30/2012 22:46:15.17 &lt;Crit:VSM.ParmInv&gt; : Argument MLAG HAL Block List has an invalid value 4982312"</pre> <pre>"07/30/2012 22:46:15.17 &lt;Crit:VSM.ParmInv&gt; : Argument Ingress Port Instance has an invalid value 16777217"</pre> <pre>"07/30/2012 22:46:15.17 &lt;Erro:Kern.Error&gt; : exvlan: handleVsmKernelRequest:7311: handleVsmKernelRequest Invalid Ingress port: 1000001 got"</pre>
xos0053941	On Summit X670 series switches, for a VPLS with two peers, the second peer counter display is incorrectly aligned in the output of command <code>show mpls statistics l2vpn</code> (second peer has extra space at start so RxBytes value is displayed near the TxPackets column).
xos0054631	On Summit X670-48x switches, ports 47 and 48 with stacking support enabled do not come up when an SFP+ transceiver (optics) module is removed, and then reinserted.
xos0055074	On Summit X670v-48x series switches, links go down after rebooting if the ports are configured with auto negotiation off and the speed is 1G on both peers when using a mini Gigabit interface converter.
xos0056125	MLAG ports move to R state after multiple executions of the command <code>restart process vsm</code> .
<b>Summit X770 Series Switches</b>	
xos0053194	AVB is not supported on the Summit X770 series switches. Time stamp support for gPTP should be added for the platform.
xos0055632	VPLS: On Summit X770 series switches, traffic is dropped after configuring LAG on service port (untagged VLAN port). The issue is with high port numbers (port 101, 102) on Summit X770 series switches.
xos0055095	Summit X770 series switches' port up/down activity can occasionally result in erroneous entries in the Layer 3 tunnel hardware table, causing Layer 3 slowpath traffic. This does not affect Layer 2 configurations, only Layer 3 configurations.
<b>ACLs</b>	
xos0056054	Disabling user-created mirror instance with ACL filter fails when a default mirror instance is enabled with ACL filter.
<b>ESRP</b>	
xos0056022	Topology change in ERPS sub-rings are not notified to the ERPS/EAPS main ring.
xos0053647	Rebooting after a failover from master to backup works fine, but after performing another failover, the ports in some ESRP member VLANs in slave are unblocked, but ports in ESRP master VLAN are blocked. This creates a situation where ports in ESRP member VLAN are unblocked in both ESRP master and slave, causing traffic loops.

**Table 16: Resolved Issues, Platform-Specific, and Feature Change Requests (CRs)**

CR Number	Description
<b>MLAG</b>	
xos0055257	After multiple reboots, VRRP instances remain in INIT state when ELSM and VRRP are running over the same single port.
xos0057384	Hardware learning is not enabled on the sharing member ports after deleting an MLAG peer.
xos0056020	L3 traffic ingressing on the ISC link from an MLAG peer is not egressing the appropriate egress MLAG port after L3 forwarding. Occurs only with jumbo frames enabled on a few of the ports.
<b>MPLS</b>	
xos0054052	Ping doesn't work on VPLS VLANs across pseudowires.
xos0055857	FDB entries are not learned after changing VMAN tag when the VMAN is associated with VPLS.
xos0055866	Multiple master switch fabric module (MSM) fail-overs cause traffic to stop completely.
xos0055914	LSP-specific transmit counters incorrectly display a zero value in the output of the command <code>show mpls statistics l2vpn</code> after failover to secondary RSVP path.
xos0055991	Traffic is dropped after disabling learning on VPLS service VLAN ports with L2VPN sharing enabled.
xos0056114	Packets are not passing over VPLS when PS tag is configured in the service VLAN.
<b>OpenFlow</b>	
xos0055315	When sending 300 packets to hit the of Default_0 flow, the switch is forwarding 3 packets to the controller through packet In message and floods the rest packets out of other OpenFlow ports.
<b>SNMP</b>	
xos0055870	The command <code>show configuration</code> should reflect deletion of default SNMP communities.
<b>VLANs</b>	
xos0056253	FDB entries are learned despite learning being disabled on ports when port specific tag added to the VLAN.
<b>VRRP</b>	
xos0055515	VRRP backup node is moving to master state despite connectivity between master and backup nodes.





# 4 ExtremeXOS Documentation Corrections

---

This chapter lists corrections to the *ExtremeXOS 15.6 User Guide*.

This chapter contains the following sections:

- [ACLs on page 110](#)
- [ACL Egress Counters Limitation on page 110](#)
- [ACL Policy Redirect on page 111](#)
- [ACL Ports Limits on page 111](#)
- [BOOTPreplay Not Supported in Virtual Routing and Forwarding \(VRF\) on page 112](#)
- [Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines on page 113](#)
- [Configure IP-MTU VLAN Command Syntax Description on page 114](#)
- [Configure Sys-Recovery-level Slot Command Platform Availability on page 115](#)
- [Debounce Commands on page 116](#)
- [ELRP and QoS on page 118](#)
- [ELRP on VPLS on page 119](#)
- [L2VPN Sharing Commands on page 120](#)
- [Link Aggregation \(LAG\) Limit for Multiprotocol Label Switching \(MPLS\) Terminated Packets on page 121](#)
- [Link Layer Discovery Protocol \(LLDP\) on page 121](#)
- [MLAG on page 122](#)
- [MLAG PIM-SM \\*,G Forwarding Limitation on page 123](#)
- [NetLogin Limitation on page 123](#)
- [NetLogin Local Authentication on page 124](#)
- [PoE Power Delivery on page 122](#)
- [PVLAN + EAPS and PVLAN + STP Recommendations on page 125](#)
- [Rate Limiting/Meters on page 126](#)
- [Remote Mirroring on page 127](#)
- [Routing Policies on page 128](#)
- [Synchronize Command on page 129](#)
- [TACACS Server on page 129](#)
- [Unconfigure Switch Erase Command on page 131](#)
- [VRRP Guidelines on page 132](#)

---

## ACLs

---

*Basic Switch Operation ExtremeXOS User Guide*, Chapter 3: “Managing the Switch”

xos0057249

The following text should be removed from multiple places under the indicated chapter:

- Only source-address match is supported.
- Access-lists that are associated with one or more applications cannot be directly deleted. They must be unconfigured from the application first, and then deleted from the CLI.
- Default counter support is added only for ACL rules and not for policy files. For policy files, you must configure count action.

*Policies and Security ExtremeXOS User Guide*, Chapter 5: “ACLs” > “ACL Rule Syntax Details”

xos0058670

Change the match conditions fields “IGMP-type number” and “IGMP-code number” to “ICMP-type number” and “ICMP-code number”.

The corresponding description fields state the correct match conditions (for example, “ICMP-type number” and “ICMP-code-number”), but the match condition fields are misprinted as “IGMP-type number” and “IGMP-code number”, respectively.

---

## ACL Egress Counters Limitation

---

*ExtremeXOS User Guide*, under *ACL Rule Syntax > Counting Packets and Bytes*

xos0061118

Add the following note:



### NOTE

Each packet increments only one counter in the egress direction. When there are multiple ACLs with action “count” applied in the port, only a single counter based on the slice priority works.

---



---

## 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

---

*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.)”

Need 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.)”

## BOOTPrelay Not Supported in Virtual Routing and Forwarding (VRF)

---

*ExtremeXOS Command Reference* and *ExtremeXOS User Guide* for DHCP/BOOTP  
xos0061332

### ExtremeXOS User Guide

Under *IPv4 Unicast Routing > DHCP/BOOTP Relay > Managing DHCP/BOOTP Relay*, add the following note::

**NOTE**

BOOTPrelay is not supported in VRF.

---

### ExtremeXOS Command Reference

For the `configure bootprelay delete` command under the *Default* heading:

Change the following from:

“If you do not specify a VR or VRF, the current VR context is used.”

To:

“If you do not specify a VR, the current VR context is used.”

---

## Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines

---

*ExtremeXOS Command Reference* for the `configure access-list vlan-acl-precedence` command

xos0060123

Change usage guidelines from:

“The following feature applies to only policy-file based ACLs that are applied on a VLAN. Use this command to switch between two VLAN-based ACL configuration modes. In the shared `vlan-aclprecedence` mode, VLAN-based ACL rules share the same precedence with other types of ACL rules. This is the default mode and provides the same behavior as in the previous software releases. In the dedicated `vlan-acl-precedence` mode, VLAN-based ACL rules have different precedence compared to other types of ACL rules. The dedicated mode yields improved installation performance for VLAN based access-lists but may affect hardware rule utilization in some configurations.”

To:

“The following feature applies to only policy-file based ACLs that are applied on a VLAN. Use this command to switch between two VLAN-based ACL configuration modes. In the shared `vlan-aclprecedence` mode, VLAN-based ACL rules share the same precedence with other types of ACL rules and provides the same behavior as in the previous software releases. In the dedicated `vlan-acl-precedence` mode, VLAN-based ACL rules have different precedence compared to other types of ACL rules and this is the default mode. The dedicated mode yields improved installation performance for VLAN based access-lists but may affect hardware rule utilization in some configurations.”

---

## Configure IP-MTU VLAN Command Syntax Description

---

*ExtremeXOS Command Reference* and *ExtremeXOS User Guide* for the `configure ip-mtu` command  
xos0061010

### Command Reference

In the Syntax Description table, change the description from:

“mtu - Specifies the IP maximum transmission unit (MTU) value. Range is from 1500 to 9194.”

To:

“mtu - Specifies the IP maximum transmission unit (MTU) value. Range is from 1,500 to 9,194. However, the command allows the maximum limit up to 9,216 considering port configuration, such as tagging which influences the L2 header size. However, values greater than 9,194 may lead to packet loss and are not recommended.”

### User Guide

Under the title *Jumbo Frames > IP Fragmentation with Jumbo Frames*:

Need to change the following content from:

“The ip-mtu value ranges between 1500 and 9194, with 1500 the default.”

To:

The ip-mtu value ranges between 1,500 and 9,194, with 1,500 the default. However, the command allows the maximum limit up to 9,216 considering port configuration, such as tagging which influences the L2 header size. However, values greater than 9,194 may lead to packet loss and are not recommended.”

---

## Configure Sys-Recovery-level Slot Command Platform Availability

---

*ExtremeXOS Command Reference* 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.”

## Debounce Commands

---

*ExtremeXOS Command Reference*

xos0060723

The following two debounce commands should appear:

### Configure stack-ports debounce time

```
configure stack-ports {port-list} debounce time [default|time]
```

#### *Description*

Configures debounce time feature on the stacking ports.

#### *Syntax Description*

port-list      Specifies one or more stacking ports.

default        Configure the default value "0"

<milliseconds> Time in milliseconds. Range is 0 (no debouncing) to 5000.

#### *Default*

Default debounce time value is 0

#### *Usage Guidelines*

Debounce timer can be configured to override the false link flaps i.e. link flaps that happens in a milliseconds interval.

#### *Example*

```
configure stack-ports 1:1 1:2 debounce time 150
```

#### *History:*

Available from ExtremeXOS 15.3.4

#### *Platforms Availability*

All stackable switches.

## Show stack-ports debounce

```
show stack-ports {port-list} debounce
```

### Description

Displays the current debounce time configured in stack-ports

### Syntax Description

port-list      Specifies one or more stacking ports.

### Default

N/A

### Usage Guidelines

To view the current debounce time configured in stack-ports. Specifying the stack-port allows to view the debounce time for particular stack-port alone.

### Example

```
show stack-ports 1:1 1:2 debounce
```

Following is the example output:

```
Stack  Debounce
Port   Time (ms)
-----
1:1    0
1:2    0
```

### History

Available from ExtremeXOS 15.3.4

### Platform Availability

All stackable switches.

## 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.

---



## ELRP on VPLS

*ExtremeXOS User Guide*, section "Using ELRP to Perform Loop Tests"

xos0057320

The following known limitation of ELRP on VPLS service VLANs for Summit X480 series switches should appear:

"On Summit X480 series switches, ELRP does not detect a loop when enabled on a VPLS service VLAN. This is a hardware limitation.

You can work around this limitation using an ACL that copies the ELRP packets to the switch:

- 1 Find out the switch's MAC address and ELRP destination MAC address.

The ELRP PDU's destination MAC address would be the switch MAC address with "01" for the first octet. For example, if the switch MAC address is "00:04:96:51:12:32", then the ELRP PDU's destination MAC address is "01:04:96:51:12:32".

- 2 Create an ACL that moves the packets to the CPU that are destined to the ELRP destination MAC address and having Self MAC as the source address.

```
create access-list elrp_lift "ethernet-source-address
<switch_ethernet_source_address>; ethernet-destination-address
<ELRP_Dest>" "copy-cpu-and-drop"
```

In this example,

```
create access-list elrp_lift "ethernet-source-address
00:04:96:51:12:32; ethernet-destination-address
01:04:96:51:12:32" "copy-cpu-and-drop"
```

- 3 Now associate the access-list with the VPLS service VLAN on which the ELRP is to be enabled or apply it to the entire switch by the use of any option.

```
configure access-list add "elrp_lift" first any
```

or

```
configure access-list add "elrp_lift" first vlan <vlan_name>
```



### NOTE

While this procedure deals with this limitation, you use one more ACL rule. So, if there are other Extreme Network devices in the service VLAN network that do not run VPLS, then it is recommended that you enable ELRP on those devices instead of using this workaround which will consume ACL resources.

## L2VPN Sharing Commands

---

*ExtremeXOS Command Reference, 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 Aggregation (LAG) Limit for Multiprotocol Label Switching (MPLS) Terminated Packets

---

In *ExtremeXOS User Guide* in the topic *Load-Sharing Algorithms > Link Aggregation Algorithms*

xos0061631

Add the following note at the end of this subtopic content:



### NOTE

In Platforms such as the Summit X670, X670v, X480, X460, and BlackDiamond 8900 series I/O modules, load sharing based on inner L3 fields in PLS-terminated packets are not supported, and the packets are forwarded as per L2 hashing.

---

## Link Layer Discovery Protocol (LLDP)

---

In the *ExtremeXOS User Guide*, under Chapter *LLDP > 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 *ExtremeXOS Command Reference Guide* for the following commands,

```
enable lldp ports
```

```
disable lldp ports
```

The existing text:

“Default: Disabled”

Should change to:

“Default: Enabled”

---

## 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.

---

## MLAG

---

*Basic Switch Operation ExtremeXOS User Guide*, under *Basic Switch Operation > MLAG > MLAG-LACP*

xos0059921

Add the following note:



### NOTE

When LACP shared ports are configured as MLAG ports, a LAG ID change after MLAG peer reboot may result in MLAG ports being removed and re-added to the aggregator. To avoid the MLAG port flap, it is recommended to configure a common LACP MAC in both the MLAG peers using the command `configure mlag peer <peer_name> lacp-mac <lacp_mac_address>`.

---

## MLAG PIM-SM \*,G Forwarding Limitation

---

*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 *NetLogin > Configuring Network Login > Exclusions and Limitations*

xos0062943

Add the following additional limitation:

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

## NetLogin Local Authentication

---

xos0063091

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

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"

## PVLAN + EAPS and PVLAN + STP Recommendations

---

xos0063369

In the *ExtremeXOS User Guide*, under chapter *VLANS > Private VLANs > PVLAN Overview > PVLAN Limitations*

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.”

## Rate Limiting/Meters

---

*ExtremeXOS User Guide* for the configure ports qosprofile command

xos0057795

Need to include the following line above the example section:

"If max-burst-size has configured as "0", then it will use maximum available burst value."

Also, change the following:

"The max-burst-size parameter is the amount of traffic above the value in the cir-rate parameter that is allowed to burst from the port(s) for a short duration."

To:

"The max-burst-size parameter is the amount of traffic above the value in the cir-rate parameter that is allowed to burst from the port(s) for a short duration. If max-burst-size has configured as "0", then it uses the maximum available burst value."



## Remote Mirroring

---

*ExtremeXOS User Guide*

xos0058665

Add information about remote mirroring guidelines:

“One-to-many remote mirroring does not work as expected where ‘mirror-to’ ports can receive double-tagged packets. This is due to hardware limitation and applies to the following platforms:

- Summit X150, X250e, X350, X450, X450e, X450a.
- BlackDiamond modules: G48T, G48P, 10G4X, G24X, a-series, e-series, c-series (except 8900 cards), and 8500 series modules

## Routing Policies

*ExtremeXOS User Guide* under *Routing Policies > Routing Policy File Syntax > Policy Action Statements*

xos0060766

In the Policy Actions table, for the "community set" attribute replace the existing text with the following text:

In the Action column:

```
"community set [no-advertise | no-export | noexport-76subconfed |
<community_num> | <as_num> : <community_num>];"
```

In the corresponding Description column:

"Replaces the existing community attribute of a route by the community specified by the action statement. Community must be enclosed in double quotes ("")."

Also, add the following note:



### NOTE

Multiple communities cannot generally be used in "community set" attribute in a BGP policy file. However, you can effectively set multiple communities by using two sets of attributes as shown in following example:

```
entry permit-anything-else {
    if {
    } then {
        community set "2342:6788";
        community add "2342:6789 2342:6790";
    }
    permit;
}
```

## Synchronize Command

*ExtremeXOS Command Reference* for the `synchronize` command

xos0059976

The following text:

“ExtremeXOS software does not allow a synchronize operation on a SummitStack between a Summit X460 or X670 switch and a Summit X480 switch. If one is attempted, the following message is displayed:...”

Should be:

“ExtremeXOS software does not allow a synchronize operation on a SummitStack between a Summit X460 ,X670 or X440 switch and a Summit X480 switch. If one is attempted, the following message is displayed:...”

## TACACS Server

*ExtremeXOS User Guide* under *Security > Authenticating Management Sessions Through a TACACS+ Server > Configuring the TACACS+ Client for Authentication and Authorization*

xos0060212

The following new topic should appear, *Changing the TACACS+ Server*:

To change a TACACS+ server configuration to avoid service interruption with respect to authentication and authorization:



### NOTE

When only a single TACACS+ server is configured, you must disable TACACS-authorization (if enabled) before reconfiguring the TACACS+ server.

- 1 Unconfigure existing primary TACACS+ server (the TACACS+ server will failover to the secondary server) by issuing the following command:

```
unconfigure tacacs server [primary | secondary]
```

- 2 Configure new primary TACACS+ server by issuing the following command:

```
configure tacacs [primary | secondary] server [ipaddress | hostname] {tcp_port} client-ip ipaddress {vr vr_name}
```

- 3 Configure the shared-secret password for the primary TACACS+ server by issuing the following command:

```
configure tacacs [primary | secondary] shared-secret  
{encrypted} string
```

**NOTE**

Only after configuring the shared-secret password for the primary server, TACACS+ will fallback to primary server from secondary.

- 4 Unconfigure the existing secondary TACACS+ server by issuing the following command:

```
unconfigure tacacs server [primary | secondary]
```

- 5 Configure the new secondary TACACS+ server by issuing the following command:

```
configure tacacs [primary | secondary] server [ipaddress |  
hostname] {tcp_port} client-ip ipaddress {vr vr_name}
```

- 6 Configure the shared-secret password for the secondary TACACS+ server by issuing the following command:

```
configure tacacs [primary | secondary] shared-secret  
{encrypted} string
```

**NOTE**

The command `disable tacacs` is not required while changing TACACS+ servers, and it is recommended to “disable tacacs-authorization” (if enabled), before disabling TACACS+.

---

## Unconfigure Switch Erase Command

---

*ExtremeXOS Command Reference* for the `unconfigure switch` command

xos0059832

Need to include the following information on the `unconfigure switch` command to explain `unconfigure switch erase`. Replace the content from the beginning of information on the command to syntax description with the following content.

“`unconfigure switch`

```
unconfigure switch {all | erase [all | nvram]}
```

Description

Returns the switch configuration to its factory default settings and reboots the switch.

Syntax Description

`all` - Specifies that the entire configuration should be changed to the default values, including the management IP address, failsafe account, and

SummitStack-specific parameters, and the switch rebooted.

`erase all` - All data such as loaded exos images(both partition), configuration files, policy files, non-volatile memory content and switch settings will be overwritten. This will render the switch inoperable until a bootrom rescue is performed. The system will reboot after the erase operation is complete which will take around 10 minutes.

`erase nvram` - Data in non-volatile memory such as selected configuration, selection image partition, log messages will be overwritten. Switch will boot up with primary image. Any unsaved configuration changes will be lost and the switch will reboot.”

## VRRP Guidelines

---

*ExtremeXOS Concepts Guide*

*Chapter 31: "VRRP" under the heading "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.”

