



ExtremeXOS Release Notes

Software Version ExtremeXOS 15.7.3-Patch1-6

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/

Support

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

Contact

Extreme Networks, Inc.
145 Rio Robles
San Jose, CA 95134
Tel: +1 408-579-2800
Toll-free: +1 888-257-3000

Table of Contents

Overview	7
New and Corrected Features in ExtremeXOS 15.7	7
Bridge and Q-Bridge Management Information Bases (MIBs)	9
Supported Platform	9
Limitations	9
Open Shortest Path First (OSPF)v2/v3 Import Policy Improvements	10
Supported Platform	10
Affected Commands	10
ExtremeXOS Trivial File Transfer Protocol (TFTP) Client MTU/Block-Size Configuration	11
Supported Platform	11
Changed CLI Commands	11
ExtremeXOS I/O Native Applications Support	12
Supported Platforms	12
New CLI Commands	12
ExtremeXOS Web-Based GUI: Chalet	13
Supported Platforms	13
Label-Switched Path (LSP) Fast Failover	13
Supported Platforms	13
Limitations	14
Changed CLI Commands	14
Multi-protocol Label Switching (MPLS) Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Explicit Route Option (ERO) Exclude Option	14
Supported Platforms	14
Changed CLI Commands	15
Clear Internet Protocol Address Resolution Protocol (IPARP)/Clear Neighbor-Discovery Refresh	17
Supported Platforms	18
Changed CLI Commands	19
Minimum Number of Link Aggregation Control Protocol (LACP) Link Aggregate Group (LAG) Members	19
Supported Platforms	20
New CLI Commands	20
Changed CLI Commands	20
AAA: Ability to Administer All Default User Accounts	21
Supported Platforms	21
New CLI Commands	21
Network Time Protocol (NTP) over Management Port	22
Supported Platforms	22
New CLI Commands	22
Limitations	22
Changed CLI Commands	22
Dynamic Host Configuration Protocol (DHCPv6) Relay—Install Routes for Snooped Delegated IPv6 Prefixes	23
Supported Platforms	23
Limitations	23
New CLI Commands	23

Multicast Listener Discovery (MLD) Source Specific Multicast (SSM) Mapping	24
Supported Platforms	24
Limitations	24
New CLI Commands	24
Changed CLI Commands	24
Virtual Machine (VM) Tracking Enhancements for NetSight	25
Supported Platforms	25
Network Logon Enhancements for NetSight	26
Supported Platforms	26
Configuring Power Budget Capability for Summit X430-8p Switches	27
CLI Commands	27
Zero Touch Provisioning	28
Supported Platforms	28
New CLI Commands	28
31-Bit Prefixes on IPv4 Interfaces (RFC 3021)	29
Supported Platforms	30
Limitations	30
Priority Flow Control (PFC) Statistics per Port	30
Supported Platforms	30
New CLI Commands	30
Entity Management Information Base (MIB) Port Support	31
Supported Platforms	32
Limitations	32
Spanning Tree Protocol (STP) Enhancements	32
Supported Platforms	32
New CLI Commands	33
Changed CLI Commands	34
Private Virtual Local Area Networks (PVLANS) Management Information Base (MIB)	35
Supported Platforms	35
MAC Address Locking	35
Supported Platforms	35
Limitations	36
New CLI Commands	36
Open Shortest Path First (OSPF)v3 Point-to-Point Interfaces	37
Supported Platforms	37
CLI Commands	38
ExtremeXOS Secure Shell/Secure Copy Protocol (SSH/SCP) Client Upgrade Using OpenSSH	
38	
Supported Platforms	38
Limitations	39
Changed CLI Commands	39
Link Aggregation Group (LAG)—Multiple VLAN Registration Protocol (MVRP) Enhancements	
39	
Supported Platforms	40
Limitations	40
Changed CLI Commands	41
OpenFlow v1.3 and Pseudowire Multiprotocol Label Switching (MPLS)	43
Supported Platforms	43
Limitations	44

Static Generalized Precision Time Protocol (gPTP) Port Roles	44
Supported Platforms	44
New CLI Commands	44
Protocol Independent Multicast (PIM) Enhancement: Shortest-Path Tree (SPT) Threshold Op- tion “Infinity”	45
Supported Platforms	45
Changed CLI Commands	45
Deprecation of Domain Field from Open Shortest Path First (OSPF)v3 Commands	46
ExtremeXOS Images for Summit X480 Series Switches	46
New Hardware Supported in ExtremeXOS 15.7	47
Hardware No Longer Supported	47
Hardware Issues in ExtremeXOS 15.6 and Later	47
Joint Interoperability Test Command (JITC) Compliance	48
ExtremeXOS Hardware/Software Compatibility and Recommendation Matrices	48
Compatibility with NetSight	48
Upgrading to ExtremeXOS	49
Downloading Supported MIBs	49
Tested Third-Party Products	49
Tested RADIUS Servers	49
Tested Third-Party Clients	50
PoE Capable VoIP Phones	50
Extreme Switch Security Assessment	51
DoS Attack Assessment	51
ICMP Attack Assessment	51
Port Scan Assessment	51
Service Notifications	51
Limits.....	52
Open Issues, Known Behaviors, and Resolved Issues	106
Open Issues	106
Known Behaviors	110
Resolved Issues in ExtremeXOS 15.7.3-Patch1-6	111
Resolved Issues in ExtremeXOS 15.7.3-Patch1-3	113
Resolved Issues in ExtremeXOS 15.7.3	117
Resolved Issues in ExtremeXOS 15.7.2	121
Resolved Issues in ExtremeXOS 15.7.1	129
ExtremeXOS Documentation Corrections.....	139
ACLs	141
ACL Egress Counters Limitation	141
ACL Policy Redirect	142
ACL Ports Limits	142
BOOTPrelay Not Supported in Virtual Routing and Forwarding (VRF)	143
ExtremeXOS User Guide	143
ExtremeXOS Command Reference	143
Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines	144

Configure IP-MTU VLAN Command Syntax Description	145
Command Reference	145
User Guide	145
Configure Sys-Recovery-level Slot Command Platform Availability	146
Debounce Commands	147
Configure stack-ports debounce time	147
Description	147
Syntax Description	147
Default	147
Usage Guidelines	147
Example	147
History:	147
Platforms Availability	147
Show stack-ports debounce	148
Description	148
Syntax Description	148
Default	148
Usage Guidelines	148
Example	148
History	148
Platform Availability	148
ELRP	149
ELRP and QoS	149
Extreme Networks Virtualization (XNV), Identity Management (IDM), and Network Time Protocol (NTP)	150
L2VPN Sharing Commands	151
Hardware Table Hash Algorithm	152
Link Aggregation (LAG) Limit for Multiprotocol Label Switching (MPLS) Terminated Packets ..	152
Link Layer Discovery Protocol (LLDP)	153
MIB	153
MLAG	154
MLAG PIM-SM *,G Forwarding Limitation	154
NetLogin Limitation	155
NetLogin Local Authentication	155
PoE Power Delivery	156
PVLAN + EAPS and PVLAN + STP Recommendations	157
Rate Limiting/Meters	158
Remote Mirroring	159
Routing Policies	160
Synchronize Command	161
TACACS Server	161
VRRP Guidelines	163

1 Overview

These release notes document ExtremeXOS® 15.7.3-Patch1-6 which resolves software deficiencies.

This chapter contains the following sections:

- [New and Corrected Features in ExtremeXOS 15.7 on page 7](#)
- [Deprecation of Domain Field from Open Shortest Path First \(OSPF\)v3 Commands on page 46](#)
- [ExtremeXOS Images for Summit X480 Series Switches on page 46](#)
- [New Hardware Supported in ExtremeXOS 15.7 on page 47](#)
- [Hardware No Longer Supported on page 47](#)
- [Hardware Issues in ExtremeXOS 15.6 and Later on page 47](#)
- [Joint Interoperability Test Command \(JITC\) Compliance on page 48](#)
- [ExtremeXOS Hardware/Software Compatibility and Recommendation Matrices on page 48](#)
- [Compatibility with NetSight on page 48](#)
- [Upgrading to ExtremeXOS on page 49](#)
- [Downloading Supported MIBs on page 49](#)
- [Tested Third-Party Products on page 49](#)
- [Extreme Switch Security Assessment on page 51](#)
- [Service Notifications on page 51](#)

New and Corrected Features in ExtremeXOS 15.7

This section lists the new and corrected features supported in the ExtremeXOS 15.7 software:

- [Bridge and Q-Bridge Management Information Bases \(MIBs\) on page 9](#)
- [Open Shortest Path First \(OSPF\)v2/v3 Import Policy Improvements on page 10](#)
- [ExtremeXOS Trivial File Transfer Protocol \(TFTP\) Client MTU/Block-Size Configuration on page 11](#)
- [ExtremeXOS I/O Native Applications Support on page 12](#)
- [ExtremeXOS Web-Based GUI: Chalet on page 13](#)
- [Label-Switched Path \(LSP\) Fast Failover on page 13](#)

- Multi-protocol Label Switching (MPLS) Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Explicit Route Option (ERO) Exclude Option on page 14
- Clear Internet Protocol Address Resolution Protocol (IPARP)/Clear Neighbor-Discovery Refresh on page 17
- Minimum Number of Link Aggregation Control Protocol (LACP) Link Aggregate Group (LAG) Members on page 19
- AAA: Ability to Administer All Default User Accounts on page 21
- Network Time Protocol (NTP) over Management Port on page 22
- Dynamic Host Configuration Protocol (DHCPv6) Relay—Install Routes for Snooped Delegated IPv6 Prefixes on page 23
- Multicast Listener Discovery (MLD) Source Specific Multicast (SSM) Mapping on page 24
- Virtual Machine (VM) Tracking Enhancements for NetSight on page 25
- Network Logon Enhancements for NetSight on page 26
- Configuring Power Budget Capability for Summit X430-8p Switches on page 27
- Zero Touch Provisioning on page 28
- 31-Bit Prefixes on IPv4 Interfaces (RFC 3021) on page 29
- Priority Flow Control (PFC) Statistics per Port on page 30
- Entity Management Information Base (MIB) Port Support on page 31
- Spanning Tree Protocol (STP) Enhancements on page 32
- Private Virtual Local Area Networks (PVLANS) Management Information Base (MIB) on page 35
- MAC Address Locking on page 35
- Open Shortest Path First (OSPF)v3 Point-to-Point Interfaces on page 37
- ExtremeXOS Secure Shell/Secure Copy Protocol (SSH/SCP) Client Upgrade Using OpenSSH on page 38
- Link Aggregation Group (LAG)—Multiple VLAN Registration Protocol (MVRP) Enhancements on page 39
- OpenFlow v1.3 and Pseudowire Multiprotocol Label Switching (MPLS) on page 43
- Static Generalized Precision Time Protocol (gPTP) Port Roles on page 44
- Protocol Independent Multicast (PIM) Enhancement: Shortest-Path Tree (SPT) Threshold Option “Infinity” on page 45

Bridge and Q-Bridge Management Information Bases (MIBs)

This feature:

- Implements the following objects/tables as defined by RFC 4363:
 - dot1qBase group
 - dot1qPortVlanTable under dot1qVlan group
 - dot1qVlanStaticTable under dot1qVlan group
- Adds dot1BasePortTable support into the existing Bridge MIB on ExtremeXOS.

Managed objects for transparent bridging are defined in the Bridge MIB. The original IEEE 802.1D is augmented by IEEE 802.1Q-2003 to provide support for virtual bridged LANs where a single bridged physical LAN network can be used to support multiple logically bridged LANs (VLANs), each of which offers a service approximately the same as that defined by IEEE 802.1D.

The Q-Bridge MIB, defined in RFC 4363 and now transferred to IEEE8021-Q-BRIDGE-MIB, provides a standard Simple Network Management Protocol (SNMP) mechanism to retrieve VLAN specific information.

Supported Platform

- 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

The groups and tables listed above are implemented as read-only.

Open Shortest Path First (OSPF)v2/v3 Import Policy Improvements

This feature allows Import Policy to be used by OSPFv2/v3 to install routes selectively into the switch routing table.

Previously, routing protocol OSPFv2/v3 applied routing policies with keyword “import-policy”, which could only be used to change the attributes of routes installed into the switch routing table. This feature now provides the flexibility of using import policy to determine the routes to be added to or removed from the routing table.

To prevent routes from being added to the routing table, the policy file must contain a matching rule with action “deny”. If there is no matching rule for a particular route, or the keyword “deny” is missing in the rule, the default action is “permit”, which means that route is installed into the routing table.

Supported Platform

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

Affected Commands

No commands have been added nor modified. The following existing commands apply import policy in OSPFv2/v3:

```
configure ospf import-policy [<policy-name> | none]
```

```
configure ospfv3 import-policy [<policy-name> | none]
```

ExtremeXOS Trivial File Transfer Protocol (TFTP) Client MTU/Block-Size Configuration

This feature adds support in ExtremeXOS commands for a configurable block-size option when transferring files using Trivial File Transfer Protocol (TFTP).

TFTP is a simple, lock-step, file transfer protocol that allows a client to get or put a file onto a remote host. TFTP is very simple to implement in a small node's limited ROM space. However, use of a 512-octet block-size is not efficient on LANs with a maximum transition unit (MTU) that may be 1,500 octets or greater.

This TFTP feature allows the client to negotiate with the server for a block-size more applicable to the network medium. This helps support TFTP over smaller-sized MTU WAN links and to take advantage of larger LAN MTUs, such as jumbo frames.

Supported Platform

- 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

Changed CLI Commands

The following commands have a configurable block-size option added (shown in bold). If you do not specify a block-size, it defaults to 1,400 bytes.

```
tftp [<ip-address> | <host-name>] {-v <vr_name>} {-b <block_size>} [-g | -p] [ {-l [ memorycard <local-file-memcard> | internal-memory <local-file-internal> | <local-file> ] } {-r <remote-file> } | {-r <remote-file> } {-l [ memorycard <local-file-memcard> | internal-memory <local-file-internal> | <local-file> ] } ]
```

```
tftp get [<ip-address> | <host-name>] {vr <vr_name>} {block-size <block_size>} <remote-file> { memory-card {<local-file-memcard>} | internal-memory {<local-file-internal>} | <local-file> } {force-overwrite}
```

```
tftp put [<ip-address> | <host-name>] {vr <vr_name>} {block-size <block_size>} [ memory-card <local-file-memcard> | internal-memory <local-file-internal> | <local-file> ] {<remote-file>}
```

```
upload configuration [<ipaddress> | <hostname>] <filename> {
{vr} <vr-name> } {block-size <block_size>}
```

```
upload log <ipaddress> {vr <vr_name>} {block-size
<block_size>} <filename> {messages [memory-buffer | nvram]}
{severity <severity> {only}} {chronological} {match <regex>}
```

```
upload debug [<ipaddress> | <hostname>] {{vr} <vrname>} {block-
size <block_size>}
```

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

```
download bootrom [<hostname> | <ipaddress>] <filename> {{vr}
<vrname>} {block-size <block_size>} {msm <slotid> | mm
<slotid> | slot <slot-number>}
```

ExtremeXOS I/O Native Applications Support

This feature provides means to extend the native capabilities of ExtremeXOS through Python processes, and includes:

- Python Software Developers Kit (SDK), which includes EXPY, an ExtremeXOS-enabled Python container.
- Python bindings to ExtremeXOS libraries, for use by EXPY and any other python-enabled process.

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

New CLI Commands

```
create process <name> python-module <python-module> {start
[auto | on-demand]} {vr <vr-name>} {description
```

```
<description> {<arg1> {<arg2> {<arg3> {<arg4> {<arg5>
{<arg6> {<arg7> {<arg8> {<arg9>}}}}}}}}}}
```

```
delete process <name>
```

```
show process
```

ExtremeXOS Web-Based GUI: Chalet

Chalet is a web-based user interface for setting up and viewing information about a switch. Chalet removes the need to know and remember commands in a command line environment. Viewable on desktop and mobile with a quick logon and intuitive navigation, Chalet features a quick setup mode for configuring switches quickly. Basic data surrounding port utilization, power, and Quality of Service (QoS) are available, and more advanced users can configure multiple VLANs, create Access Control Lists (ACLs), and configure Audio Video Bridging (AVB). Chalet is packaged with the ExtremeXOS 15.7.1 image for all platforms.

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

Label-Switched Path (LSP) Fast Failover

ExtremeXOS Multiprotocol Label Switching (MPLS) provides support for Redundant Resource Reservation Protocol - Traffic Engineering (RSVP-TE) label-switched paths (LSPs), which allows the configuration of a primary and up to two secondary paths for a particular LSP. To reduce the traffic loss over a Redundant RSVP-TE LSP when switching to a different path, ExtremeXOS MPLS now allows the path to be switched without dataplane updates to any routes or pseudowires that are using the LSP. This enhancement significantly reduces the number of dataplane changes and results in reduced traffic loss during the switch. Primary LSP paths using fast-reroute (FRR) detour LSPs also benefit from this enhancement when the point of local repair (PLR) is at the ingress node.

Supported Platforms

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

Limitations

- Failover times are always affected by failure detection times.
- Only redundant RSVP-TE LSPs (including primary FRR LSPs) benefit from this feature.
- Circuit Emulation Services (CES) Pseudowires (PWs) do not benefit from this feature.

Changed CLI Commands

Debug commands, such as `debug mpls show lspdb` now display the new abstracted tunnel index, and `debug mpls show tech nhlfe-index` displays the nhlfe-index usage.

Multi-protocol Label Switching (MPLS) Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Explicit Route Option (ERO) Exclude Option

This feature allow the path for a Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Label-Switched Path (LSP) to be calculated to avoid certain hops.

With traffic engineering RSVP-TE LSPs, it is often advantageous to prevent the primary and secondary paths of a Redundant RSVP-TE LSP from overlapping because if the paths overlap, the LSP may not be protected, since if that section of the path becomes non-functional, both paths will go down. The `mplsTunnelHopTable` defined in RFC 3812 allows a hop to be defined as “include” or “exclude”. ExtremeXOS previously only supported “include”, meaning that the defined hop had to be included in the path calculation. This feature adds support for “exclude”, which allows you to define a hop that must be avoided in the path calculation.

Supported Platforms

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

Changed CLI Commands

```
configure mpls rsvp-te path <path_name> add ero {{include}
<ipNetmask> [strict | loose] | exclude <ipNetmask>} {order
<number>}
```

The following show commands now show “exclude” option path information (shown in bold):

```
show mpls rsvp-te path
```

```
(pacman debug) J1.15 # show mpls rs path
```

```
Path Name #LSP #ERO Ord# ERO IP Netmask Type Inc/Exc
```

```
-----
path1      1      1    100 9.50.1.2/32 strict include
path2      1      1    100 9.50.1.2/32 n/a      exclude
```

```
show mpls rsvp-te path detail
```

```
(pacman debug) J1.16 # show mpls rs path det
```

```
Path Name : path1
```

```
  Hop List Index : 1
```

```
  Path Option Index: 1
```

```
  #ERO/Hops      : 1
```

```
  #LSP References : 1
```

```
  ERO            : Order# IpAddress/Mask Type Inc/Exc
```

```
-----
                : 100      9.50.1.2/32 strict include
```

```
Path Name : path2
```

```
  Hop List Index : 2
```

```
  Path Option Index : 1
```

```
  #ERO/Hops      : 1
```

```
  #LSP References : 1
```

```
  ERO            : Order# IpAddress/Mask Type Inc/Exc
```

: 100 9.50.1.2/32 n/a **exclude**

```
show mpls rsvp-te lsp detail
```

```
(pacman debug) J1.17 # show mpls rs lsp det
```

```
Ingress LSP Name: lsp_to_baha
```

```
Destination      : 11.100.100.4 Admin Status : Enabled
```

```
IP Traffic       : Allow #VPLS Cfgd      : 0
```

```
VPN Traffic      : Allow #VPLS In-Use   : 0
```



```

Path Name: path1

Oper Status : Enabled UpTime : 0d:0h:4m:49s
Profile Name : default
Peak Rate : 0 Kbps Max Burst Size : 0 Kb
Committed Rate : 0 Kbps Setup/Hold Priority: 7/0
Record Route : Enabled

MTU : Use Local I/F

Tunnel ID : 1 Ext Tunnel ID : 11.100.100.1
LSP ID : 0 State Changes : 1
LSP Type : Primary Bandwidth Cfgd : False
Activity : Active
Failures : 0 Retries-since last failure : 0
Retries-Total : 0

Configured ERO: Order IP Address/Mask Type Inc/Exc
100 9.50.1.2/32          strict include

Advertised Label: n/a Received Label : 0x00434
Rx Packets : n/a Tx Packets : --
Rx Bytes : n/a Tx Bytes : --
Next Hop I/F : 9.50.1.1 - j1-j2vlan1
Next Hop Addr : 9.50.1.2

Record Route : Indx  IP Address Label
                :    1    9.50.1.2  0x00434
                :    2    9.54.1.4  0x00434

```

Clear Internet Protocol Address Resolution Protocol (IPARP)/Clear Neighbor-Discovery Refresh

This feature enhances the `clear iparp` and `clear neighbor-discovery` commands by adding an ability to clear only inactive neighbor entries. When refresh keyword is specified, `clear iparp refresh` and `clear neighbor-discovery refresh` solicit all neighbor entries, delete all non-responding/

inactive entries and keep all active entries, without impacting IPv4/IPv6 traffic for active entries.

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

Changed CLI Commands

Changes are shown in bold:

```
clear iparp {<ip_addr> {{vr} <vr_name>} | {vlan}
<vlan_name> | {vr} <vr_name>} {refresh}

clear neighbor-discovery {cache {ipv6}} { <ipv6_addr>
{{vr} <vr_name>} | {vlan} <vlan_name> | {vr} <vr_name> }
{refresh}
```

Show Iparp and Show Neighbor-Discovery Output

When `clear iparp refresh` or `clear neighbor-discovery refresh` is executed, the switch sends out ARP request/neighbor solicit for every suitable entry and marks each entry with indicating that ARP response or neighbor advertisement has not been received from the neighbor. After an ARP response or neighbor advertisement is received from the neighbor, the asterisk (*) is cleared from the entry.

Minimum Number of Link Aggregation Control Protocol (LACP) Link Aggregate Group (LAG) Members

The LAG minimum links feature allows you to configure a value for the minimum number of active links to keep the entire LAG up.

For example, for a LAG consisting of four ports and the minimum links set to two, at least two links must be up for the LAG to be up. When the LAG falls below two active links, the entire LAG is brought down, and all applications using this LAG are informed that the LAG is down. Previously, the implicit minimum link value was one, meaning that if there was at least one link up, the entire LAG stayed up.

For static LAGs, the number of active physical member port links is checked to see if it is greater than or equal to the user-configured minimum link value. If so, the LAG remains up. If this test fails, the static LAG is brought down and all applications receive a Link-Down message for this LAG. As soon as the number of active physical member ports equals or exceeds the configured minimum link value, the static LAG is brought up and all applications receive a Link-Up message for this LAG.

In the case of LACP, how many member ports LACP has requested to be added to the LAG is tracked. After successfully negotiating with its peer, LACP sends a request to add a member port to the LAG. When the number of member ports LACP has requested to be added to the LAG drops below the configured minimum link value, the LACP LAG is brought down and all applications receive a Link-Down message for this LAG. As soon LACP has added enough member ports so that the total equals or exceeds the configured minimum link value, the LAG is brought up and all applications receive a Link-Up message for this LAG.

Both static and LACP LAGs can be used with Multiple Link Aggregation Group (MLAG) on either the ISC or the MLAG ports.

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

New CLI Commands

```
configure sharing <port> minimum-active <min_links_active>
```

Changed CLI Commands

The following show commands now display the configured minimum link value:

```
show port <port> sharing
```

```
show lacp lag <port> detail
```

AAA: Ability to Administer All Default User Accounts

Accounts can be disabled or enabled locally using read/write access. Even all administrative privileged accounts and user privileged accounts can be disabled. Lawful-Intercept accounts are disabled under user privileged option.

This enable/disable command affects the following North Bound Interfaces (NBIs) in management access realm:

- Console
- Telnet
- SSH
- HTTP
- XML

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

New CLI Commands

```
enable account [all {admin | user} | <name>]
```

```
disable account [all {admin | user} | <name>]
```

Network Time Protocol (NTP) over Management Port

Network Time Protocol (NTP) is a protocol for synchronizing clocks of servers or network entities using TCP/IP-based networks which have a coherent variable latency. It is designed particularly to resist the effects of variable latency by using a jitter buffer. NTP provides Coordinated Universal Time Clock (UTC). However, no information about time zones or daylight saving time is transmitted. NTP uses a hierarchical, semi-layered system of levels of clock sources. Each level of this hierarchy is termed a stratum and is assigned a layer number starting with 0 (zero) at the top. The stratum level defines its distance from the reference clock and exists to prevent cyclical dependencies in the hierarchy.

This feature adds configuring NTP on any one virtual router at a time. By default NTP is configured on VR-Default, which can be changed to other virtual routers.

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

New CLI Commands

```
configure ntp vr <vr-name>
```

Limitations

NTP must be disabled globally before changing the virtual router for NTP, and should be enabled again afterwards. All present NTP VLAN configurations are deleted when changing the virtual router.

Changed CLI Commands

The output of the show ntp command now shows which virtual router NTP has been configured on (shown in bold):

```
XNTP                : Enabled
Authentication      : Disabled
Broadcast-Client    : Disabled
VR configured      : <vr-name>
```

Dynamic Host Configuration Protocol (DHCPv6) Relay— Install Routes for Snooped Delegated IPv6 Prefixes

This feature provides a way for the ExtremeXOS DHCPv6 Relay Agent to install routes for ensuring reachability to DHCPv6 clients that get their IP addresses from customer edge switches which distribute a delegated IPv6 prefix to the clients.

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

This feature enhances the behavior of the Relay Agent, which resides between a delegating router and a requesting router, in adding routing information for the snooped delegated prefixes. It does not provide the capability to be a delegating router, which can delegate prefixes, nor a requesting router, which can allocate IPv6 addresses after receiving a delegated prefix.

New CLI Commands

```
configure bootprelay ipv6 prefix-delegation snooping [on
{vlan} <vlan_name> | off [{vlan} <vlan_name> | vlan all]]
```

```
clear bootprelay ipv6 prefix-delegation snooping [{ipv6-
prefix} <ipv6_prefix> | ipv6-prefix all] [{vlan}
<vlan_name> | vlan all]
```

```
configure bootprelay ipv6 prefix-delegation snooping add
<ipv6_prefix> <ipv6Gateway> {vlan} <vlan_name> valid-time
<valid_time>
```

```
show bootprelay ipv6 prefix-delegation snooping {{vlan}
<vlan_name>}
```

Multicast Listener Discovery (MLD) Source Specific Multicast (SSM) Mapping

This feature enables MLDv1 hosts to participate in Source Specific Multicast (SSM).

The Multicast Listener Discovery (MLD) Source Specific Multicast (SSM) Mapping feature is an IPv6 equivalent of the IPv4 feature, Internet Group Management Protocol (IGMP) SSM Mapping. The MLD SSM Mapping feature allows you to configure mapping entries, thereby enabling MLDv1 hosts to participate in SSM functionality by sending MLDv1 reports. You can configure the sources and group/group ranges for which SSM functionality has to be applied. You can also configure Domain Name System (DNS) names for a group/group range.

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

- Only 50 sources (static or dynamic) are allowed for each group address/group range. The DNS server may send only 15 IPv6 source addresses in its response, thereby limiting the number of dynamic sources supported.
- Only one DNS name is allowed for each group address/group range.

New CLI Commands

```
[enable | disable] mld ssm-map {{vr} <vrname>}

configure mld ssm-map add <v6groupnetmask> [<v6sourceip> |
<src_domain_name>] {{vr} <vrname>}

configure mld ssm-map delete <v6groupnetmask> [<v6sourceip>
| <src_domain_name> | all] {{vr} <vrname>}

unconfigure mld ssm-map {{vr} <vrname>}

show mld ssm-map {<v6groupnetmask>} {{vr} <vrname>}

refresh mld ssm-map <v6groupnetmask> {{vr} <vrname>}
```

Changed CLI Commands

Changes are in bold.

For the `refresh igmp` command “`dns group`” is now optional:

```
refresh igmp ssm-map {dns group} [<grpipaddress> <netmask>
| <ipNetmask>] {{vr} <vrname>}
```

Virtual Machine (VM) Tracking Enhancements for NetSight

Previously, when XNV is enabled on Multiple Link Aggregation Group (MLAG)-enabled ports, both MLAG peers authenticate the VM independently. Now, to integrate with NetSight and Network Access Control (NAC), provided MLAG peers have ISC connectivity, only one MLAG peer authenticates a VM that is learned on an MLAG port.

When ISC connectivity between MLAG peers is established, the peer with the highest IP address is chosen to be the authenticator. This peer authenticates a virtual machine (VM) based on the chosen authentication method. The result of the authentication is checkpointed by the authenticator to its peer so that the same virtual port profile (VPP) gets applied to the VM on both peers.

When the MLAG peer that is the authenticator goes down, the other peer detects that the authenticator is down and re-authenticates the VM at the next authentication interval. The peer that takes over as the authenticator does not re-authenticate the VMs immediately, but waits for the re-authentication timer to expire.

VMs learned on non-MLAG ports are authenticated by the detecting peer.

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

Network Logon Enhancements for NetSight

Network logon (Netlogon) controls the admission of user packets into a network by allowing MAC addresses from users that are properly authenticated. Netlogon is controlled on a per-port basis. When Netlogon is enabled on a port, that port does not forward any packets until authentication occurs. Netlogon is capable of three types of authentication: Web-based, MAC-based, and 802.1X.

This feature add support for multiple authentication protocols on a netlogon-enabled port. You must specify the authentication protocol priority or order per port, which dictates the action for the client or supplicant that is getting authenticated on this port. You can use the command line to configure the authentication protocol order. By default the protocol precedence order for a netlogon enabled port is:

- 1 Dot1x
- 2 MAC
- 3 Web-based

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

Configuring Power Budget Capability for Summit X430-8p Switches

You can now configure the budgeted Power over Ethernet (PoE) power reserved for all power devices for Summit X430-8p switches from 60–90 W.

The default power budget is 60 W. If you set the reserved budget to greater than the default value, a warning message notifies you that you can exceed the default budget only on a standalone desktop configuration.

If the power consumed by devices on the switch exceeds the configured budget, new ports requiring power are either denied or other lower priority ports are disconnected.

CLI Commands

```
configure inline-power budget <num_watts> [slot <slot> |  
NULL]
```

```
unconfigure inline-power budget [slot <slot> | NULL]
```

```
show inline-power
```

Zero Touch Provisioning

The Zero Touch Provisioning feature enables switches “out of the box” to automatically gain a management IP address and configuration without serial cables and manual configuration:

- **Management port address:** Configures an IPv4 Link-Local IP management port address allowing for easier access to the switch by connecting with a Web browser to `http://0xa9fe<last 2 MAC chars>`, where `<last 2 MAC chars>` is the last two alphanumeric groups in the MAC address found on the switch label.

For example:

For MAC address: 00:04:96:97:E9:EE

Use `http://0xa9fee9ee` (or just `0xa9fee9ee`)

- **Receives and Option 43 or Option 125 messages** (if both Option 43 and 125 messages are in the DHCP reply, Option 125 is implemented):
 - Option 43 for ExtremeXOS image updates; configuration, policy, and script files
 - Option 125 for NetSight trap address

The Zero Touch Provisioning feature starts automatically when booting up a switch, unless one the following conditions is true:

- Configuration file is present.
- Management IP address is present.
- License has been configured.
- You enter `disable auto-provisioning` command.
- debug mode is enabled.

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

New CLI Commands

```
disable auto-provision
```

```
show auto-provision {{vr} <vr_name>}
```

31-Bit Prefixes on IPv4 Interfaces (RFC 3021)

You can now configure IPv4 addresses with 31-bit prefixes on both network VLANs and the Management VLAN. Applications (for example, Ping) and protocols (for example, OSPF) can use the IPv4 interfaces configured with these 31-bit prefixes.

After the introduction of Classless Inter-Domain Routing (CIDR), the 32-bit IPv4 address has been partitioned into two parts that go by various names but can be called the prefix and the host-number:

IP-address = <Prefix> <Host-number>

Two host number values are traditionally reserved:

- All ones: This value indicates a directed broadcast on the network indicated by the prefix.
- All zeroes: With the original “classed” definition of IPv4 addresses, the address value indicated the length of the prefix, so a prefix (to be used as a route key, for example) could be determined from just an IPv4 address. To distinguish IPv4 addresses that specified prefixes from those that specified hosts, the host-number value of all-zeroes was reserved for IPv4 addresses that specified prefixes.

RFC 3021 defines the practice of not reserving the all-ones and all-zeros host number values on networks with 31-bit prefixes.

There is no new commands for this feature. The only changes are values allowed in the commands.

For example:

```
configure {vlan} <vlan_name> ipaddress [ <ipaddress>
{<netmask>} | <ipNetmask>]
```

255.255.255.254 is allowed for <netmask>

/31 is allowed in <ipNetmask> (for example, 192.168.0.1/31)

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

Directed broadcasts cannot be sent from other networks to a network using a 31-bit prefix.

Priority Flow Control (PFC) Statistics per Port

This feature introduces a new command that displays the number of priority flow control (PFC) pause control frames sent and received on a specific port.

Supported Platforms

The following platforms support the PFC feature in the new command:

- BlackDiamond 8800 switches with 8900-MSM128 and the following switch modules:
 - BlackDiamond 8900-10G24X-c modules (manufacturing number 800397-00)
 - BlackDiamond 8900-40G6X-xm modules, 40G ports and 10G ports when in 4x10 partition mode
- The following Summit series switches:
 - Summit X460 switches, 10G ports
 - Summit X460G2 switches, 1G and 10G ports
 - Summit X670 switches, 1G and 10G ports
 - Summit X670V switches, 1G, 10G, and 40G ports
 - Summit X670-G2 switches, 1G, 10G, and 40G ports
 - Summit X770-32q, 10G and 40G ports

New CLI Commands

```
show port {port_list} flow-control {rx-pauses} {tx-pauses}
{no-refresh}
```

Entity Management Information Base (MIB) Port Support

The Entity management information base (MIB) provides a standard Simple Network Management Protocol (SNMP) mechanism to retrieve device-specific information. It provides a standard mechanism to indicate device hierarchy using SNMP, to determine what physical entities are managed by the agent, and thereby to be able to communicate with the agent about a particular physical entity and understand its hierarchical position in the chassis.

ExtremeXOS 15.7 adds the following:

- Supports entPhysicalMfgDate, entPhysicalUris, and entPhysicalUUID.

Previously, the Entity MIB in ExtremeXOS was implemented based on RFC 2737. The enhancement to the latest version 4 as defined by RFC 6933 adds the following three additional MIB objects for each physical entity: entPhysicalMfgDate, which contains the manufacturing date of the managed entity; entPhysicalUris, which provides additional identification information about the physical entity; and entPhysicalUUID, which provides a unique identification about the physical entity.

- Implements entAliasMappingTable and entPhysicalContainsTable.

The enhancement to the latest version of the MIB adds three additional groups. Previously, of the various groups defined by the RFC for the Entity MIB, ExtremeXOS only implemented the entityPhysical group. The other groups defined in the RFC are entityLogical group, entityMapping group, entityGeneral group, and entityNotification group. The group entityLogical describes the logical entities managed by a single agent. In ExtremeXOS, an agent represents a single logical entity; therefore, this table has not been implemented. The group entityMapping contains three tables. The entLPMappingTable table, which contains mappings between logical entities and the physical components supporting that entity, is not implemented since ExtremeXOS only supports one logical entity. The tables entAliasMappingTable and entPhysicalContainsTable are implemented for ExtremeXOS 15.7.

- Implements physical port support for the entityPhysical group

Physical ports are currently not supported in entityPhysical Group. They will be added into entityPhysical Group, grouped by slot as appropriate using relationships specified in entPhysicalContainedIn.

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

entPhysicalUri and entPhysicalAssetID are implemented as read-only objects.

Spanning Tree Protocol (STP) Enhancements

This enhancement to the ExtremeXOS Spanning Tree Protocol (STP) implementation provides:

- Loop protect

STP requires switches to send Bridge Protocol Data Units (BPDUs) bi-directionally and build the topology database. If a port only allows traffic one-way (Uni-directional link failure), and the switch still sees that port as up, a loop can form without the switch detecting it. This loop protect feature avoids loops by checking if a BPDU timeout occurs on a port, and then changes the state to “discarding” until a BPDU is received.

- Backup root bridge

The backup root bridge becomes the root bridge when the primary root bridge is unavailable due to failure, malfunction, or administrative errors.

- Multisource detection

Multisource detection is a feature that prevents network disruption due to excessive topology changes caused by a full duplex port transmitting multiple BPDUs with different source MAC addresses, and hence different BPDU information.

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

New CLI Commands

```
clear {stpd} <stpd_name> ports <port_list> protocol-  
migration  
  
configure {stpd} <stpd_name> backup-root [on | off]  
  
configure {stpd} <stpd_name> ports auto-edge [on | off]  
<port_list>  
  
configure {stpd} <stpd_name> ports loop-protect [on | off]  
<port_list>  
  
configure {stpd} <stpd_name> ports loop-protect partner  
[capable | incapable] <port_list>  
  
configure {stpd} <stpd_name> tx-hold-count <tx_hold_count>  
  
configure {stpd} <stpd_name> ports restricted-tcn [on |  
off] <port_list>  
  
configure {stpd} <stpd_name> trap backup-root [on | off]  
  
configure {stpd} <stpd_name> trap new-root [on | off]  
  
configure {stpd} <stpd_name> trap bpdu-restrict [on | off]  
  
configure {stpd} <stpd_name> trap loop-protect [on | off]  
  
configure {stpd} <stpd_name> trap topology-change {edge-  
ports} [on | off]  
  
configure {stpd} <stpd_name> dispute-threshold [<threshold>  
| none] <port_list>  
  
configure {stpd} <stpd_name> loop-protect event-window  
<interval_time>  
  
configure stpd filter-method [system-wide| port-based]
```

Changed CLI Commands

The following commands are modified to show various aspects of this feature.

```
show stp port
```

Modified to show loop protect, partner capabilities, restricted TCN status, operator edge, and auto edge status.

```
show stp
```

Modified to show all trap status, loop protect threshold value, and loop protect event window values.

```
show stp detail
```

Modified to show all trap status loop protect, threshold value, and loop protect event window values.

```
show stp ports non forwarding reason
```

Modified to show the non-forwarding reason on ports.

```
show stp ports blocked ports
```

Modified to show all blocked ports in spanning tree.

```
show stp ports counters
```

Modified to show all counters for spanning tree.

```
show MSTP digest value
```

Modified to shows MSTP digest value.

```
configure stpd filter-method [system-wide| port-based]
```

```
show stpd
```

Modified to show STP filter method.

Private Virtual Local Area Networks (PVLANS) Management Information Base (MIB)

ExtremeXOS 15.7 now has the following MIB tables:

- PVLAN Table supporting GET/SET.
- PVLAN Member VLAN Table supporting GET/SET

and the following changes to existing MIB tables:

- Modifications to extremeVlanOpaqueControlTable for adding translated ports.
- Modifications to extremeVlanOpaqueTable to display translated VLAN ports.

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

MAC Address Locking

This feature introduces the ability to limit access to a port to specified MAC addresses or a maximum number of MAC addresses on a first-come first-served basis. MAC address locking locks a port to one or more MAC addresses, preventing connection of unauthorized devices via a port. With MAC locking enabled, the only frames forwarded on a MAC address locked port are those with the configured or dynamically selected MAC addresses for that port.

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

- MAC address locking cannot be enabled along with existing limit learning and lock learning features.
- MAC address locking is not supported on ports in the VPLS service VLAN.

New CLI Commands

```
enable mac-locking
```

```
disable mac-locking
```

```
enable mac-locking [port_list | all]
```

```
disable mac-locking [port_list | all]
```

```
configure mac-locking ports <port_list> learn-limit-action  
[disable-port | remain-enabled]
```

```
configure mac-locking ports <port_list> first-arrival  
limit-learning <learn_limit>
```

```
configure mac-locking ports <port_list> static limit-  
learning <learn_limit>
```

```
configure mac-locking ports <port_list> static [add |  
enable | disable] station <station_mac_address>
```

```
configure mac-locking ports <port_list> static delete  
station [<station_mac_address> | all]
```

```
configure mac-locking ports <port_list> first-arrival aging  
[enable | disable]
```

```
configure mac-locking ports <port_list> first-arrival move-  
to-static
```

```
configure mac-locking ports <port_list> first-arrival link-  
down-action [clear-macs | retain-macs]
```

```
configure mac-locking ports <port_list> trap {violation |  
threshold} [on | off]
```

```
configure mac-locking ports <port_list> log {violation |  
threshold} [on | off]
```

```
clear mac-locking station [all | {mac
<station_mac_address>} {first-arrival | static} {ports
<port_list>}]

clear mac-locking disabled-state ports <port_list>

show mac-locking {ports <port_list>}

show mac-locking stations {first-arrival | static} {ports
<port_list>}
```

Open Shortest Path First (OSPF)v3 Point-to-Point Interfaces

This feature for OSPFv3 is similar to OSPFv2 point-to-point interfaces, which is already supported in ExtremeXOS. Point-to-point interfaces skip designated router/backup designated router (DR/BDR) election and reach operational state faster. The links on these interfaces do not require network link-state advertisement (LSAs) and the type of intra-area prefix LSAs generated by DRs. Network topologies that can use links with only one neighbor can take advantage of the speed and efficiency of point-to-point interfaces.

Supported Platforms

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

CLI Commands

The command to configure OSPFv3 point-to-point interfaces was available in ExtremeXOS, but support was missing:

```
configure ospfv3 {domain <domainName>} add [{vlan} <vlan-name> |  
{tunnel} <tunnel-name>] {instance-id <instanceId>} area <area-  
identifier> {link-type [auto | broadcast | point-to-point]}  
{passive}
```

There are no changes to the command, but support is added to configure the link-type for an interface and the warning message that previously appeared when link-type was specified is removed:

```
Warning! Link type configuration is not supported
```

```
Link type will always be automatically determined
```

ExtremeXOS Secure Shell/Secure Copy Protocol (SSH/SCP) Client Upgrade Using OpenSSH

This feature upgrades SSH client/SCP client behavior using OpenSSH from openssh-3.9p1 to openssh-6.5p1.

Only SSH version 2 is supported.

SCP file transfers to the switch include the following file types: configuration (cfg), script (xsf and py), policy (pol), SSH key (ssh).

SCP file transfers from the switch include the following file types: configuration (cfg), script (xsf and py), policy (pol), and SSH key (ssh).

Supported Platforms

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

Limitations

- Only password-based authentication is supported for SSH/SCP client.
- SCP client does not support upload of image or BootROM files (for example: xos, xmod, and xtr).
- SCP client does not support download of image or BootROM files (for example: xos, xmod, and xtr) from EXOS SCP server.
- Only supports transfer of files such as cfg, xsf, py, pol, and ssh files to/from the switch.
- Current version of openssl is not Federal Information Processing Standards (FIPS) compliant.

Changed CLI Commands

The following commands are changed for this feature. Removed items are struck through and added items are in bold.

```
ssh2 {cipher [3des | blowfish] <cipher>} {mac <mac>}
{compression [on | off]} {port <port>} {user <username>} {vr
<vr_name>} <user@host> {<remote_command>}
```

```
cp2 {cipher [3des | blowfish] <cipher>} {mac <mac>}
{compression [on | off]} {port <port>} {vr <vr_name>} [
<user@host:file> <local-file> | <local-file> <user@host:file> ]
```

Link Aggregation Group (LAG)—Multiple VLAN Registration Protocol (MVRP) Enhancements

Link aggregation allows an increased bandwidth and resilience by using a group of ports to carry traffic in parallel between switches. Multiple ports can be aggregated into one logical port. MVRP can be enabled on the logical port. The MVRP control packets are transmitted on any available physical port of the LAG. The peer on the other side receives the packet and processes it as if it is received on the logical port. MVRP supports both dynamic (LACP) as well as static load sharing. MVRP data structure is based on port Instance. All dynamic VLANs created or propagated for a given port are stored for each port Instance. For normal ports, the port Instance corresponds to the PIF port instance, and for LAG ports, the port Instance corresponds to the LIF port Instance. The port instance is not shown in any of the standard show commands, though it is available as a part of the debug commands. Once MVRP is enabled on the master port, addition/deletion of individual links is supported. MVRP packets received on the newly added link are accounted instantaneously.

Supported Platforms

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

Limitations

- The individual ports of the LAG, including the master port, should not have MVRP configuration prior to grouping.
- MVRP can be enabled/disabled only on master port. The individual links cannot be configured.
- Once sharing is disabled, MVRP configuration of the master port is lost (default is disabled).
- The statistics and counters shown on the MVRP show commands are a cumulative counter for all links added together. There are no per link counters.
- The actual load sharing of the traffic is beyond MVRP's domain and should take place as per the configured LAG setting. MVRP just adds the LAG port to the VLAN(s).

Changed CLI Commands

The following command is updated to show MVRP information. Changes are shown in bold.

```
show vlan <vlan_name>
```

```
show vlan sys_vlan_0100
```

```
VLAN Interface with name sys_vlan_0100 created dynamically by MVRP
```

```

      Admin State:      Enabled          Tagging:
802.1Q Tag 100

      Description:      None
      Virtual router:   VR-Default
      IPv4 Forwarding:  Disabled

IPv6 Forwarding: Disabled

      IPv6:              None
      STPD:              None
      Protocol:          Match all unfiltered protocols
      Loopback:          Disabled
      NetLogin:          Disabled
      QosProfile:        None configured
      Egress Rate Limit Designated Port: None configured
      Flood Rate Limit QosProfile: None configured
      Ports:            1. (Number of active ports=1)
                        Tag: *2gH

Flags: (*) Active, (!) Disabled, (g) Load Sharing port
      (b) Port blocked on the vlan, (m) Mac-Based port
      (a) Egress traffic allowed for NetLogin
      (u) Egress traffic unallowed for NetLogin
      (t) Translate VLAN tag for Private-VLAN
      (s) Private-VLAN System Port, (L) Loopback port

```

- (e) Private-VLAN End Point Port
- (x) VMAN Tag Translated port
- (G) Multi-switch LAG Group port
- (H) Dynamically added by MVRP**

OpenFlow v1.3 and Pseudowire Multiprotocol Label Switching (MPLS)

This feature provides limited support of OpenFlow v1.3 and includes controlling a Multiprotocol Label Switching (MPLS) Pseudowire overlay network using OpenFlow:

- Adds support for OpenFlow version 1.3 by upgrading to OpenVswitch (OVS) version 2.1.
- The OpenFlow Group table is supported for MPLS flows only; otherwise, only a single table is supported.
- L2VPN with static Pseudowire and static MPLS tunnel (using OpenFlow vendor extensions).
- MPLS label switching (for Pseudowire only).
- Forwarding between normal VLAN/VMAN and OpenFlow Pseudowires using standard Ethernet flooding and learning.
- Enabling OpenFlow-controlled traffic and normal traffic to be sent and received on the same VLAN/port.
- Adds a new global “hybrid” mode operation to allow switches to operate as a normal switch, but also allow a controller to install flows that can affect any port and/or any VLAN.
- Continues to support existing OpenFlow capabilities and scale that were implemented in previous version of ExtremeXOS (v15.4.1).

Supported Platforms



NOTE

MPLS features are only supported on platforms that support MPLS. The Summit X430 and X440 do not support MPLS.

- BlackDiamond X8 [all modules; single Master Switch Fabric Module (MSM) only]
- BlackDiamond 8800 [8900 (XL-series) and C-series; single Management Module (MM) only]
- Summit X770, X670, X480, X460, and X440

Limitations

- This features is not a full implementation of OpenFlow (see earlier list of supported OpenFlow v1.3 capabilities)
- Does not include implementing a controller or application.
- MPLS and Pseudowire instances are limited by platform capabilities.
- Access controll list (ACL)-based rule scale is limited by platform capabilities.
- Forwarding database (FDB)-based rules are limited by platform FDB sizes.
- Failover is not supported on chassis.
- Stacking is not supported.

Static Generalized Precision Time Protocol (gPTP) Port Roles

The Generalized Precision Time Protocol (gPTP) port role feature allows you to enable or disable the Best Master Clock Algorithm (BMCA) function of gPTP. Enabling BMCA selects the Grandmaster, builds the time-synchronization spanning tree, and selects the port roles. If you disable BMCA and the switch should not be the Grandmaster Clock, then you must select the ports on which gPTP is enabled to take on the SlavePort role. If no port is configured for the SlavePort role, the switch behaves as the Grandmaster, and all ports on the which gPTP is enabled take on the MasterPort role.

Supported Platforms

Summit X770, X670, X670-G2, X460, X460-G2, and X440, and X430 series switches

New CLI Commands

```
configure network-clock gptp bmca [on | off]
```

```
configure network-clock gptp slave-port [<port_no> | none]
```

Protocol Independent Multicast (PIM) Enhancement: Shortest-Path Tree (SPT) Threshold Option “Infinity”

This enhancement adds a new option for the shortest-path tree (SPT) threshold parameter that causes the last hop router (LHR) to not switch over to SPT from rendezvous point (RP) tree. When this option is configured, an LHR or intermediary route does not build (S;G)s entries; instead traffic paths are based on (*;G)s only.

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 following command is changed (shown in bold):

```
configure {ipv4|ipv6} pim spt_threshold infinity
```

Deprecation of Domain Field from Open Shortest Path First (OSPF)v3 Commands

For ExtremeXOS 15.7, the `domain` field is deprecated from all OSPFv3 commands.

ExtremeXOS Images for Summit X480 Series Switches

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

Table 1: Summit X480 Series Switches Software Image Files

	Main Install image	Diagnostic image
Content	All Summit X480 content (except diagnostics)	Summit X480 diagnostics
File Name	<code>summitX480-15.6.xx.yy.xos</code>	<code>summitX480-15.6.xx.yy-diagnostics.xmod</code>
File Type	Standard ExtremeXOS image	XMOD image
Installation Notes	<ul style="list-style-type: none"> Installing the main SummitX480 image over a previous release leaves the previous installation of the diagnostics image intact, as it is stored separately from the main ExtremeXOS image. You can continue to use the previously installed diagnostic version to run diagnostics. The Summit XMODs, such as SSH can be used with the summitX480 ExtremeXOS image. 	To update to a newer version of the diagnostics, you download and install the latest XMOD version. The diagnostics XMOD can be installed to the active or standby partition and diagnostics can be used immediately. There is no need to reboot or any other action to complete the installation.

The following scenarios will produce an error or warning message:

- Not having the diagnostic image installed on a Summit X480 series switch or slot.
- Installing the main Summit X480 image without the diagnostics image present.
- Installing the general Summit image (`summitX-15.7.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-15.7.xx.yy.xos` file image, and this image installs the diagnostics capability.

New Hardware Supported in ExtremeXOS 15.7

This section lists the new hardware supported in ExtremeXOS 15.7:

- BDXB-100G4X-XL I/O modules for the BlackDiamond X8 series switches

Hardware No Longer Supported

The following hardware is no longer supported in ExtremeXOS 15.7.1:

- 8500-G24X-e
- 8500-G48T-e
- 8500-MSM24

Hardware Issues in ExtremeXOS 15.6 and Later

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/Software Compatibility and Recommendation Matrices

The *ExtremeXOS Hardware/Software Compatibility and Recommendation Matrices* 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/Software Compatibility and Recommendation Matrices* can be found at:

www.extremenetworks.com/documentation

Compatibility with NetSight

ExtremeXOS 15.7 is compatible with NetSight version 6.1 and later.

Upgrading to ExtremeXOS

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

- Beginning with ExtremeXOS 12.1, an ExtremeXOS core image (.xos file) must be downloaded and installed on the alternate (non-active) partition. If you try to download to an active partition, the error message **"Error: Image can only be installed to the non-active partition."** 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
- For Summit X480 series switches, starting with ExtremeXOS 15.6, two separate software image files are used for both individual switches and stacks that include Summit X480 series switches. For more information, see [ExtremeXOS Images for Summit X480 Series Switches](#) on page 46.

Downloading Supported MIBs

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

<https://esupport.extremenetworks.com/>

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 15.7.3-Patch1-6.

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
AAA (local) —maximum number of admin and local user accounts.	All platforms	8
Access lists (meters) —maximum number of meters.	BlackDiamond 8000 series e-series, group of 24 ports c-series 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 and BDX X8-100G4X-XL 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
Access lists (policies) —suggested maximum number of lines in a single policy file.	All platforms	300,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Access lists (policies) —maximum number of rules in a single policy file. ^a	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 BlackDiamond XB-100G4X-XL modules	8,192 ingress, 1,024 egress 139,264 ingress 1,024 egress.
	Summit X440, X430 group of 24 ports	1,024 ingress
	Summit X460, E4G-400	
	Summit X480	4,096 ingress, 512 egress (up to) 61,440 ingress, 1,024 egress
	Summit X670 VIM4-40G4x	2,048 ingress 1,024 egress
	Summit X480	
	Summit X480 VIM3-40G4X	8,192 ingress/ 1,024 egress 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
Access lists (slices) —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 8900-40G6X-xm	17 ^b 10 ingress, 4 egress
	BlackDiamond X8 a-series modules	10 ingress, 4 egress
	BlackDiamond X8-100G4X modules	16 ingress, 4 egress
	BlackDiamond XB-100G4X-XL modules	17 ingress
	E4G-200	4 egress 8 ingress, 4 egress
	Summit X440, X430	4 ingress
	Summit X460, E4G-400, X460-G2	
	Summit X480	16 ingress, 4 egress
	Summit X670 VIM4-40G4x	17 ^b ingress, 4 egress 10 ingress, 4 egress
	Summit X480 VIM3-40G4X Summit X770, X670-G2	10 ingress, 4 egress 12 ingress 4 egress
	AVB (audio video bridging) —maximum number of active streams. NOTE: It is recommended that you do not use on more than 8 ports on this switch.	Summit X440, X460, X460-G2
Summit X670, X670, X670-G2		4,096
Summit X430		100*

Table 2: Supported Limits (Continued)

Metric	Product	Limit
BFD sessions —maximum number of BFD sessions.	All platforms (default timers—1 sec)	512
	BlackDiamond X8 and 8800 (minimal timers—50 msec)	10 ^c
	All Summits (minimal timers—100 msec)	10 ^c
BGP (aggregates) —maximum number of BGP aggregates.	All platforms (except E4G-200, X430, and X440) with Core license or higher	256
BGP (networks) —maximum number of BGP networks.	All platforms (except E4G-200, X430, and X440) with Core license or higher	1,024
	BlackDiamond X8 series	1,024
BGP (peers) —maximum number of BGP peers. NOTE: With default keepalive and hold timers.	BlackDiamond X8 series	512
	BlackDiamond 8000 series	512
	BlackDiamond xl-series	512
	All Summits, except X480, X440, X430, E4G-200	128*
	E4G-400	128*
	Summit X480	512
BGP (peer groups) —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
BGP (policy entries) —maximum number of BGP policy entries per route policy.	All platforms (except E4G-200, X430, and X440) with Core license or higher	256
BGP (policy statements) —maximum number of BGP policy statements per route policy.	All platforms (except E4G-200, X430, and X440) with Core license or higher	1,024
BGP multicast address-family routes —maximum number of multicast address-family routes.	BlackDiamond 8900 xl-series	524,256 (up to) ^b
		25,000
	BlackDiamond X8 series	1,048,544 (up to) ^m
	BlackDiamond X8-100G4X-XL modules	25,000
	Summit X460, X460-G2, X670, X670-G2, X770	25,000
	Summit X480	524,256 (up to) ^b
E4G-400	25,000	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
BGP (unicast address-family routes) —maximum number of unicast address-family routes.	BlackDiamond 8900 xl-series	524,256 (up to) ^b 25,000
	BlackDiamond X8 series	1,048,544 (up to) ^m
	BlackDiamond X8-100G4X-XL modules	25,000
	Summit X460, X460-G2, X670, X670-G2, X770	524,256 (up to) ^b
	Summit X480 E4G-400	25,000
BGP (non-unique routes) —maximum number of non-unique BGP routes.	BlackDiamond 8900 xl-series	1,200,000
	BlackDiamond X8 series	24,000
	BlackDiamond X8 with 100G4X-XL modules	1,200,000
	Summit X460, X460-G2, X670, X670-G2, X770	25,000
	Summit X480 E4G-400	1,000,000 25,000
BGP ECMP —maximum number of equalcost multipath for BGP and BGPv6.	All platforms, except Summit X430, X440, and E4G-200	2, 4, or 8
	BlackDiamond 8800 G48Te2 (for BGPv6)	N/A
BGPv6 (unicast address-family routes) —maximum number of unicast address family routes.	BlackDiamond 8900 xl-series	20,000
	BlackDiamond 8800 c-series	6,000
	BlackDiamond 8000 e-series	240
	BlackDiamond X8 series	8,000
	BlackDiamond X8-100G4X-XL modules	20,000
	Summit X460, X460-G2	6,000
	Summit X480	20,000
	Summit X670, X670-G2, X770 E4G-400	8,000 6,000
BGPv6 (non-unique routes) —maximum number of non-unique BGP routes	BlackDiamond 8900 xl-series	24,000
	BlackDiamond 8800 c-series	18,000
	BlackDiamond 8000 e-series	720
	BlackDiamond X8 series	24,000
	BlackDiamond X8-100G4X-XL modules	24,000
	Summit X460, X460-G2	18,000
	Summit X480, X670, X670-G2, X770	24,000
	E4G-400	18,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per virtual router.	All platforms, except Summit X430	4
BOOTP/DHCP relay —maximum number of BOOTP or DHCP servers per VLAN.	All platforms, except Summit X430	4
CES TDM pseudowires —maximum number of CES TDM pseudowires per switch.	E4G-200 and E4G-400	256
Connectivity fault management (CFM) —maximum number of CFM domains. NOTE: With Advanced Edge license or higher.	All platforms	8
CFM —maximum number of CFM associations. NOTE: With Advanced Edge license or higher.	All platforms	256
CFM —maximum number of CFM up end points. NOTE: With Advanced Edge license or higher.	BlackDiamond 8000 series BlackDiamond X8 series Summit series	32 32 32
CFM —maximum number of CFM down end points. NOTE: With Advanced Edge license or higher.	BlackDiamond 8000 series BlackDiamond X8 series Summit series X460, E4G-200, E4G-400 (non-load shared ports) All other platforms	32 32 256 (non-load shared ports) 32 (load shared ports) 32
CFM —maximum number of CFM remote end points per up/down end point. NOTE: With Advanced Edge license or higher.	All platforms	2,000
CFM —maximum number of dot1ag ports. NOTE: With Advanced Edge license or higher.	All Summits, except X430	128
CFM —maximum number of CFM segments. NOTE: With Advanced Edge license or higher.	All platforms	1,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
CFM —maximum number of MIPs. NOTE: With Advanced Edge license or higher.	All platforms	256
CLEAR-Flow —total number of rules supported. The ACL rules plus CLEAR-Flow rules must be less than the total number of supported ACLs.	BlackDiamond X8, BlackDiamond 8800 Summit X440, X430 Summit X670 Summit X460, X460-G2, X770, X670-G2 Summit X480 E4G-200 E4G-400	4,096 1,024 2,048 4,094 8,192 2,048 4,094
Data Center Bridging eXchange (DCBX) protocol Type Length Value (TLVs) —maximum number of DCBX application TLVs.	All platforms	8
DHCPv6 Prefix Delegation Snooping —Maximum number of DHCPv6 prefix delegation snooped entries.	All platforms	256 (with Underlying Protocol Ripng) 128 (with Underlying protocol OSPFv3) 1,024 (with static routes)
Dynamic ACLs —maximum number of ACLs processed per second. NOTE: Limits are load dependent.	Summit X480, X670 with 50 DACLs with 500 DACLs BlackDiamond X8 BlackDiamond 8800	10 5 N/A N/A
EAPS domains —maximum number of EAPS domains. NOTE: An EAPS ring that is being spatially reused cannot have more than four configured EAPS domains.	BlackDiamond 8000 series BlackDiamond X8 series Summit series (except X430), E4G-200, E4G-400 Summit X430	64 64 32 4
EAPsv1 protected VLANs —maximum number of protected VLANs.	BlackDiamond 8000 series BlackDiamond X8 series Summit series, E4G-200, E4G-400	2,000 2,000 1,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
EAPSV2 protected VLANs —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
ELSM (vlan-ports) —maximum number of VLAN ports.	BlackDiamond 8000 series	5,000
	BlackDiamond X8 series	5,000
	All Summits, E4G-200, E4G-400	5,000
ERPS domains —maximum number of ERPS domains without CFM configured	BlackDiamond 8800 series	32
	BlackDiamond X8 series	32
	Summit series (except X430), E4G-200, E4G-400	32
	Summit X430	4
ERPS domains —maximum number of ERPS domains with CFM configured.	BlackDiamond 8800 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
ERPSv1 protected VLANs —maximum number of protected VLANs.	BlackDiamond 8800 series	2,000
	BlackDiamond X8 series	2,000
	All Summits, E4G-200, E4G-400	1,000
ERPSv2 protected VLANs —maximum number of protected VLANs	BlackDiamond 8800 series	2,000
	BlackDiamond X8 series	2,000
	All Summits (except X430), E4G-200, E4G-400	500
ESRP groups —maximum number of ESRP groups.	All platforms	7
ESRP domains —maximum number of ESRP domains.	All platforms	64
ESRP VLANs —maximum number of ESRP VLANs.	BlackDiamond 8800	1,000
	BlackDiamond X8	2,048
	All Summits	1,000
	E4G-200, E4G-400	1,000
ESRP (maximum ping tracks) —maximum number of ping tracks per VLAN.	All platforms (except Summit X430)	8
ESRP (IP route tracks) —maximum IP route tracks per VLAN.	All platforms (except Summit X430)	8

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
FDB (blackhole entries) — maximum number of multicast 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
FDB (maximum L2 entries) — maximum number of MAC addresses.	BlackDiamond 8000 c-series	32,768 ^f
	BlackDiamond 8000 e-series	8,192 ^f
	BlackDiamond 8000 (system), except 8900 xl-series	128,000 ^f
	BlackDiamond 8900 xl-series	524,488 (up to) ^b
	BlackDiamond X8 a-series modules	128,000 ^f
	BlackDiamond X8-100G4X modules	384,000 ^f
	BlackDiamond BDx X8-100G4X-XL modules	1,048,576 (up to) ^{b g}
	E4G-200, E4G-400	32,000 ^f
	Summit X440, X430	16,000 ^f
	Summit X480 (40G4X)	524,488 (up to) ^b 32,000 ^f
	Summit X460	288,000 ^f
	Summit X460-G2, X670-G2	128,000 ^f
	Summit X670	294,912 ^{e f}
	Summit X770	
FDB (Maximum L2 entries) — maximum number of multicast 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
FIP Snooping VLANs	BlackDiamond X8	768
	BlackDiamond 8800 (8900-40G6X-c only)	
	Summit X670	
FIP Snooping Virtual Links (FPMA mode) per port group	BlackDiamond X8	1,908
	BlackDiamond 8800 (8900-40G6X-c only)	
	Summit X670	

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Identity management —recommended number of ACL entries per identity. NOTE: Number of ACLs per identity based on system ACL limitation.	All platforms, except Summit X430. Summit X430	20 N/A
Identity management —maximum number of dynamic ACL entries configured as an individual dynamic rule, or as an ACL entry in a policy file.	All platforms, except Summit X430. Summit X430	500 N/A
IGMP sender —maximum number of IGMP senders per switch (IP multicast compression disabled). ^r NOTE: 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 and BDX X8-100G4X-XL L modules E4G-200, E4G-400 Summit X440 Summit X480 Summit X460 Summit X460-G2 Summit X670 Summit X670-G2 Summit X770 Summit X430	2,048 ^h 500 ⁱ 2,048 ^h 4,096 ^h 3,000 ⁱ 4,096 ^h 4,096 ⁱ 16,384 ^e 2,048 64 4,096 2,048 4,096 3,000 ⁱ 4,096 ⁱ 4,096 64

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IGMP sender —maximum number of IGMP senders per switch (IP multicast compression enabled). NOTE: Assumes source-group-vlan mode. For additional limits, see: <ul style="list-style-type: none"> • Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode. on page 79 • Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode. on page 79 	BlackDiamond 8800 c-series	2,048 ⁱ
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900-10G24X-c modules	2,048 ⁱ
	BlackDiamond 8900-G96T-c modules	4,096 ⁱ
	BlackDiamond 8900-40G6X-xm	3,000 ⁱ
	BlackDiamond 8900 xl-series	12,000 ⁱ
	BlackDiamond X8 a-series modules	4,096 ^b
	BlackDiamond X8-100G4X and BDx X8-100G4X-XL modules	64,000 ⁱ
	E4G-200	3,000 ^{i,j}
	E4G-400	6,000 ^{i,j}
	Summit X440	192 ⁱ
	Summit X460	6,000 ⁱ
	Summit X460-G2	21,000 ⁱ
	Summit X480	12,000 ⁱ
	Summit X670	3,000
Summit X770, X670-G2	66,500 ⁱ	
Summit X430	192	
IGMP snooping per VLAN filters —maximum number of VLANs supported in per-VLAN IGMP snooping mode.	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 and BDx X8-100G4X-XL 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, X670-G2	2,000	
IGMPv1/v2 SSM-map entries —maximum number of IGMPv1/v2 SSM mapping entries.	All platforms	500
IGMPv1/v2 SSM-MAP entries —maximum number of sources per group in IGMPv1/v2 SSM mapping entries.	All platforms	50

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per port. ^s	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
IGMPv2 subscriber —maximum number of IGMPv2 subscribers per switch. ^s	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
IGMPv3 maximum source per group —maximum number of source addresses per group.	All platforms	250
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per port. ^s	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	4,000
	Summit X460, E4G-400	2,000
IGMPv3 subscriber —maximum number of IGMPv3 subscribers per switch. ^s	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, X460-G2, E4G-400	20,000
	Summit X770, X670-G2	30,000
IP ARP entries in software —maximum number of IP ARP entries in software. NOTE: May be limited by hardware capacity of FDB (maximum L2 entries).	BlackDiamond X8-100G4X modules	229,374 (up to) ^k
	Summit X670-G2, X770	131,072(up to) ^k
	Summit X670, X480, X460, X440, X430	20,480
	Summit X460-G2	57,344 (up to) ^k
	E4G-200, E4G-400	20,480

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IP ARP entries in software with distributed mode on —maximum number of IP ARP entries in software with distributed mode on.	BlackDiamond 8000 series with 8900-MSM128 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
IPv4 ARP entries in hardware with distributed mode on —maximum number of IP ARP entries in hardware with distributed mode on	Per BlackDiamond 8900-10G8X-xl, up to 260,000 per system	32,500 ^b
	Per BlackDiamond 8900-G48X-xl or 8900-G48T-xl, up to 130,000 per system	16,250 ^b
	Per BlackDiamond 8000 c-series, up to 18,000 per system	8,000
	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	
IPv4 ARP entries in hardware with minimum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with minimum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series switches, assumes number of IP route reserved entries is 100 or less.	BlackDiamond 8800 c-, xm-series	8,000
	BlackDiamond 8000 e-series	1,000 ⁱ
	BlackDiamond 8900 xl-series	16,000
	BlackDiamond X8 a-series	16,000
	BlackDiamond X8-100G4X modules	182,000(up to) ^k ^l
	BlackDiamond BDX X8-100G4X-XL modules	282,000 ^l
	E4G-200	8,000
	E4G-400	16,000
	Summit X440	412
	Summit X670, X480 (40G4X)	8,000
	Summit X460, X480	16,000
	Summit X460-G2	50,000 (up to) ^k
Summit X770, X670-G2	108,000(up to) ^k	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IPv4 ARP entries in hardware with maximum LPM routes —maximum recommended number of IPv4 ARP entries in hardware, with maximum LPM routes present. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is “maximum.”	BlackDiamond 8800 c-, xm-series	6,000 ⁱ
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900 xl-series	12,000 ⁱ
	BlackDiamond X8 a-series	12,000 ⁱ
	BlackDiamond X8-100G4X modules	172,000 (up to) ^k 1
	E4G-200	6,000 ⁱ
	E4G-400	12,000 ⁱ
	Summit X440	380
	Summit X460, X480	12,000 ⁱ
	Summit X670, X480 VIM3-40G4X	6,000 ⁱ
Summit X770, X670-G2	98,000 (up to) ^k	
Summit X460-G2	43,000 (up to) ^k	
IP flow information export (IPFIX) —number of simultaneous flows.	BlackDiamond 8900 xl-series modules	4,096 ingress, 4,096 egress
	BlackDiamond 8900 c-series modules	4,096 ingress, 4,096 egress
	BlackDiamond X-100G4X and BDX X8-100G4X-XL 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	2,048 ingress, 2,048 egress
	E4G-400	4,096 ingress, 4,096 egress 2,048 ingress, 2,048 egress

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IPv4 remote hosts in hardware with zero LPM routes —maximum recommended number of IPv4 remote hosts (hosts reachable through a gateway) in hardware when LPM routing is not used. For BlackDiamond 8800, BlackDiamond X8, E4G, and Summit series, assumes number of IP route reserved entries is 0, and number of IPv4 ARP entries present is 100 or less.	BlackDiamond 8800 c-series	18,000 ⁱ
	BlackDiamond 8000 e-series	1,000 ⁱ
	BlackDiamond 8900 xl-series	40,000 ^b
	BlackDiamond 8900-40G6X-xm	22,000 ⁱ
	BlackDiamond X8 a-series	28,000 ⁱ
	BlackDiamond X8-100G4X and X8-100G4X-XL modules	311,000 (up to) ^k l
	E4G-200	18,000 ⁱ
	E4G-400	20,000 ⁱ
	Summit X440	448
	Summit X460	20,000 ⁱ
	Summit X460-G2	73,000 ^k
	Summit X480	40,000 ^b
	Summit X670, X480 VIM3-40G4X	22,000 ⁱ
Summit X770, X670-G2	176,000 (up to) ^k	
IPv4 routes —maximum number of IPv4 routes in software (combination of unicast and multicast routes).	BlackDiamond 8900 xl-series with 8900-MSM128 or MSM-48c	524,256 (up to) ^b 25,000
	All other BlackDiamond 8000 series hardware	25,000
	BlackDiamond X8 series	1,048,544 (up to) ^m
	BlackDiamond X8 with BDX X8-100G4X-XL modules	256
	Summit X440	25,000
	Summit X460, X670, X770, X670-G2, X460-G2	524,256 (up to) ^b
	Summit X480	25,000
	E4G-200, E4G-400	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IPv4 routes (LPM entries in hardware) — number of IPv4 routes in hardware.	BlackDiamond 8800 c-series	12,000
	BlackDiamond 8000 e-series	480
	BlackDiamond 8900 xl-series	524,256 (up to) ^b _n
	BlackDiamond 8900-40G6X-xm	16,000 ^e
	BlackDiamond X8 series	16,000 ^e
	BlackDiamond BDX X8-100G4X-XL modules	1,048,544 (up to) ^o
	E4G-200, E4G-400	12,000
	Summit X440	32
	Summit X460, X460-G2	12,000
	Summit X480	524,256 (up to) ^b _n
	Summit X480 VIM3-40G4X	16,000 ⁿ
	Summit X670	12,000
Summit X770, X670-G2	16,000	
IPv6 addresses on an interface — maximum number of IPv6 addresses on an interface.	All platforms	255
IPv6 addresses on a switch — maximum number of IPv6 addresses on a switch	BlackDiamond 8000 series	512
	BlackDiamond X8 series	2,048
	E4G-200, E4G-400	512
	Summit X440	254
	Summit X460, X480	512
	Summit X770, X670, X670-G2, X460-G2	2,048

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IPv6 host entries in hardware —maximum number of IPv6 neighbor entries in hardware.	BlackDiamond 8800 c-, xm-series	3,000 ⁱ
	BlackDiamond 8000 e-series	250 ⁱ
	BlackDiamond 8900-10G24X-c modules	2,000 ⁱ
	BlackDiamond 8900-G96T-c modules	4,000 ⁱ
	BlackDiamond 8900 xl-series	8,192 (up to) ^{b i}
	BlackDiamond X8 a-series	3,000 ⁱ
	BlackDiamond X8-100G4X and BDx X8-100G4X-XL modules	49,000 ^{i p}
	E4G-200	2,000 ⁱ
	E4G-400	3,000 ⁱ
	Summit X440	192 ⁱ
	Summit X460, X670, X480 VIM3-40G4X	3,000 ⁱ
	Summit X770, X670-G2	36,750 ⁱ
	Summit X480, X670v-48t	6,000 ⁱ
	Summit X460-G2	22,000 ⁱ
IPv6 routes (LPM entries in hardware) —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) ^b
	BlackDiamond X8 series	8,000
	BlackDiamond BDx X8-100G4X-XL modules	524,288 (up to) ^p
	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) ^b
IPv6 routes with a mask greater than 64 bits in hardware —maximum number of such IPv6 LPM routes in hardware.	BlackDiamond 8000 c-, e-, xm-series	256
	BlackDiamond 8000 xl-series	245,760 (up to) ^b
	BlackDiamond X8 series	256
	BlackDiamond BDx X8-100G4X-XL modules	524,288 (up to) ^p
	E4G-200, E4G-400	256
	Summit X440, X460, X460-G2, X670, X670-G2, X770, X480 (VIM3-40G4X)	256
	Summit X480	245,760 (up to) ^b

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IPv6 route sharing in hardware —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)	0-128
	Summit X460-G2, X670-G2, X770	0-64 (> 64 single path only)
	BlackDiamond X8 a-series	0-128
	BlackDiamond X8-100G4X modules	0-64 (> 64 single path only)
	BlackDiamond X8-100G4X-XL modules	0-128 ^P
IPv6 routes in software —maximum number of IPv6 routes in software.	Summit X440, X430	N/A
	BlackDiamond 8800 G48Te2	N/A
	BlackDiamond 8900 xl-series with 8900-MSM128 or MSM-48c	245,760 (up to) ^b
	All other BlackDiamond 8000 series hardware	25,000
	BlackDiamond X8 series	25,000
	BlackDiamond X8 with BDX X8-100G4X-XL modules	524,288 (up to) ^P
	Summit X460, X460-G2, X670, X670-G2, X770, E4G-200, E4G-400	25,000
Summit X480	245,760 (up to) ^b	
Summit X440	256	
IP router interfaces —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
	BlackDiamond 8800	512
	Summit X440	254
	Summit X480, X460	512
	E4G-200, E4G-400	512
IP multicast static routes —maximum number of permanent multicast IP routes.	All platforms (except Summit X430, X440)	1,024
	Summit X430, X440	32
IP unicast static routes —maximum number of permanent IP unicast routes.	All platforms (except Summit X430, X440)	1,024
	Summit X430, X440	32

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IP route sharing (maximum gateways) —Configurable maximum number of gateways used by equal cost multipath OSPF, BGP, IS-IS, static routes, or L2VPNs. Routing protocol OSPF is limited to 16 ECMP gateways per destination. Routing protocols BGP and IS-IS are limited to 8 ECMP gateways per destination. Static routes are limited to 32 next-hops. L2VPNs are limited to 16 LSPs per pseudowire on platforms that support 32 gateways, and 64 LSPs per pseudowire on platforms that support 64 gateways.	All platforms, except Summit X430, X440, X670, and BlackDiamond X8	2, 4, 8, 16, or 32
	Summit X670, BlackDiamond X8	2, 4, 6, 8, 16, 32, or 64
	Summit X430, X440	N/A
	BlackDiamond 8800 G48Te2 (for IPv6)	N/A
IP route sharing (total destinations) —maximum number of unique destinations used by multipath OSPF, OSPFv3, BGP, IS-IS, or static routes. NOTE: For platforms with limit of 524,256 or higher, the total number of "destination+gateway" pairs is limited to 2,097,024. For example, if the number of unique destinations is 524,256, only 2 gateways per destination is supported. 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.	BlackDiamond 8800 c-series	12,256
	BlackDiamond 8000 e-series	480
	BlackDiamond 8900 xl-series	524,256 (up to) ^b
	BlackDiamond 8900-40G6X-xm	16,352
	BlackDiamond X8	16,352
	BlackDiamond BDX X8-100G4X-XL modules	1,048,544 (up to) ^m
	E4G-200, E4G-400	12,256
	Summit X480	524,256 (up to) ^b
	Summit X670, X670-G2, X770, X480 (VIM3-40G4X)	16,352
	Summit X460-G2, X460	12,256

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IP route sharing (total combinations of gateway sets) —maximum number of combinations of sets of adjacent gateways used by multipath OSPF, BGP, IS-IS, or static routes.	BlackDiamond 8800 c-, xl-, and xm-series	
	default maximum gateways of 4	510
	if maximum gateways is 2	1,022
	if maximum gateways is 8	254
	if maximum gateways is 16	126
	if maximum gateways is 32	62
	BlackDiamond 8000 e-series	
	default maximum gateways of 4	30
	if maximum gateways is 2	62
	if maximum gateways is 8	14
	if maximum gateways is 16	6
	if maximum gateways is 32	2
	BlackDiamond X8 series, Summit X670	
	default maximum gateways of 4	
	if maximum gateways is 2	510
	if maximum gateways is 8	1,022
	if maximum gateways is 16	254
	if maximum gateways is 32	126
	if maximum gateways is 64	62
	if maximum gateways is 64	30
	Summit X460, X460-G2, X480, X670, X670-G2, X770, E4G-200, E4G-400	
	default maximum gateways of 4	
	if maximum gateways is 2	510
	if maximum gateways is 8	1,022
	if maximum gateways is 16	254
	if maximum gateways is 32	126
	if maximum gateways is 32	62
IP multinetting (secondary IP addresses) —maximum number of secondary IP addresses per VLAN.	BlackDiamond 8800	64
	BlackDiamond X8	64
	All Summits, except X440, X430	255
	Summit X440	32
IS-IS adjacencies —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	256
	E4G-400	128
IS-IS ECMP —maximum number of equal cost multipath for IS-IS.	All platforms, except Summit X440, X430	2, 4, or 8
	BlackDiamond 8800 G48Te2 (for IPv6)	N/A
IS-IS interfaces —maximum number of interfaces that can support IS-IS.	All platforms, except Summit X440, x430	255

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IS-IS routers in an area —recommended maximum number of IS-IS routers in an area.	Summit X480	128
	All other platforms, except Summit X440, X430	256
IS-IS route origination —recommended maximum number of routes that can be originated by an IS-IS node.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond BDX X8-100G4X-XL modules	30,000
	BlackDiamond 8900 xl-series	30,000
	Summit X460, X460-G2, X670, X670-G2, X770, X480,	20,000
	E4G-400 E4G-200	20,000 25,000
IS-IS IPv4 L1 routes in an L1 router —recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	BlackDiamond 8000 series	25,000
	BlackDiamond X8 series	25,000
	BlackDiamond BDX X8-100G4X-XL modules	120,000
	BlackDiamond 8900 xl-series	120,000
	Summit X480	50,000
	Summit X460, X460-G2, X670, X670-G2, X770	25,000
IS-IS IPv4 L2 routes —recommended maximum number of IS-IS Level 2 routes.	E4G-200, E4G-400	25,000
	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	25,000
	BlackDiamond BDX X8-100G4X-XL modules	120,000
	BlackDiamond 8900 xl-series	120,000
	Summit X480	50,000
IS-IS IPv4 L1 routes in an L1/L2 router —recommended maximum number of IS-IS Level 1 routes in an L1/L2 IS-IS router.	Summit X460, X460-G2, X670, X670-G2, X770	25,000
	E4G-200, E4G-400	25,000
	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	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IS-IS IPv6 L1 routes in an L1 router —recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond BDX X8-100G4X-XL modules	40,000
	BlackDiamond 8900 xl-series	40,000
	Summit X480	25,000
	Summit X460, X460-G2, X670, X670-G2, X770, E4G-400	10,000
IS-IS IPv6 L2 routes —recommended maximum number of IS-IS Level 2 routes.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond X8 BDX X8-100G4X-XL modules	40,000
	BlackDiamond 8900 xl-series	40,000
	Summit X480	15,000
	Summit X460, X460-G2, X670, X670-G2, X770	10,000
IS-IS IPv6 L1 routes in an L1/L2 router —recommended maximum number of IS-IS Level 1 routes in a L1/L2 router.	E4G-200, E4G-400	10,000
	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond BDX X8-100G4X-XL modules	15,000
	BlackDiamond 8900 xl-series	15,000
	Summit X480	15,000
IS-IS IPv4/IPv6 L1 routes in an L1 router —recommended maximum number of IS-IS Level 1 routes in a Level 1 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	Summit X460, X460-G2, X670, X670-G2, X770, E4G-400	10,000
	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond BDX X8-100G4X-XL modules	60,000
	BlackDiamond 8900 xl-series	60,000
	Summit X480	40,000
E4G-200, E4G-400	Summit X460, X460-G2, X670, X670-G2, X770	20,000
		20,000
		20,000
		20,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
IS-IS IPv4/IPv6 L2 routes in an L2 router —recommended maximum number of IS-IS Level 2 routes in a Level 2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	20,000
	BlackDiamond BDX X8-100G4X-XL modules	60,000
	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
IS-IS IPv4/IPv6 L1 routes in an L1/L2 router —recommended maximum number of IS-IS Level 1 routes in a Level 1/Level2 IS-IS router. The numbers documented are based on 50% IPv4 routes and 50% IPv6 routes.	BlackDiamond 8000 series	20,000
	BlackDiamond X8 series	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
Jumbo frames —maximum size supported for jumbo frames, including the CRC.	All platforms	9,216
L2 VPN: VCCV (pseudowire Virtual Circuit Connectivity Verification) VPNs per switch —maximum number of VCCV enabled VPLS VPNs.	All platforms, except Summit X440, X430	16
L2 VPN: VPLS MAC addresses —maximum number of MAC addresses learned by a switch.	BlackDiamond 8900 xl-series	512,000
	BlackDiamond 8900-40G6X-xm	128,000
	BlackDiamond X8 a-series modules	128,000
	BlackDiamond X8-100G4X modules	384,000
	BlackDiamond BDX X8-100G4X-XL modules	1,048,576 ^g
	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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
L2 VPN: VPLS VPNs —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
L2 VPN: VPLS peers —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
L2 VPN: LDP pseudowires —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	
L2 VPN: static pseudowires —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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
L2 VPN: Virtual Private Wire Service (VPWS) VPNs —maximum number of virtual private networks per switch.	Summit X460	1,000
	Summit X480, X770	4,000
	Summit X480-40G VIM	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	
Layer-2 IPMC forwarding caches —(IGMP/MLD/PIM snooping) in mac-vlan mode. NOTE: IPv6 and IPv4 L2 IPMC scaling is the same for this mode.	BlackDiamond 8800 e-series switches	2,000
	BlackDiamond 8800 c- and xl-series switches	8,000
	BlackDiamond 8800 xm-series switches	15,000
	BlackDiamond X8 series switches	15,000
	E4G-200, E4G-400	15,000
	Summit X480, X460	8,000
	Summit X670, X670V	8,000
	Summit X440	15,000
	Summit X770, X670-G2	4,000
	Summit X460-G2	77,500 ^k
Summit X430	24,576 ^q	
5,000		
Layer-2 IPMC forwarding caches —(IGMP/MLD/PIM snooping) in mixed-mode. NOTE: IPv6 and IPv4 L2 IPMC scaling is the same for this mode.	BlackDiamond 8800 e-series switches	N/A
	BlackDiamond 8800 xl- and c-series switches	8,000
	BlackDiamond 8800 xm-series switches	15,000
	BlackDiamond X8, Summit X670, X670V	15,000
	E4G-200 and E4G-400 cell site routers, Summit X460	8,000
	Summit X440	4,000
	Summit X770, X670-G2	77,500 ^k
	Summit X460-G2	24,576 ^q
	Summit X480	8,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Layer-3 IPMC forwarding caches —(PIM, MVR, PVLAN) in mixed-mode. ¹ NOTE: IPv6 L3 IPMC scaling is 50% of these limits in this mode.	BlackDiamond 8800 e-series switches	N/A
	BlackDiamond 8800 xl- and c-series switches	6,000
	BlackDiamond 8800 xm-series switches	3,000
	BlackDiamond X8 a-series modules	6,000
	BlackDiamond X8-100G4X and modules	64,000
	E4G-200 cell site routers, Summit X670	3,000
	E4G-400 cell site routers, Summit X460, X480, X670V	6,000
	Summit X440	192
	Summit X770, X670-G2	77,500 ^k
	Summit X460-G2	21,000 ^k
Load sharing —maximum number of loadsharing groups. NOTE: The actual number of load-sharing groups that can be configured is limited by the number of physical ports present in the switch or SummitStack.	BlackDiamond 8000 series without 8900-40G6X-xm	
	With distributed IP ARP mode off (default)	128
	With distributed IP ARP mode on	64
	BlackDiamond 8000 series 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
Load sharing —maximum number of ports per load-sharing group. NOTE: For custom algorithm *For L2 and L3 algorithms NOTE: For a mix of Summit X770 and Summit X670 series switches in a stack, the limits are the Summit X670 limits.	BlackDiamond X8 series	64
	Summit X460-G2 (standalone)	32
	Summit X670 (standalone)	32 * 16 **
	Summit X670 (stacked) Summit X670-G2 (stacked)	64 * 16 **
	Summit X770 (standalone) Summit X670-G2 (standalone) Summit X460-G2 (standalone)	32
	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
Logged messages —maximum number of messages logged locally on the system.	All platforms	20,000
MAC address learning rate —hardware learning rate	E4G-200	22 msec
MAC-based security —maximum number of MAC-based security policies.	All platforms	1,024
MAC Locking —Maximum number of MAC locking stations that can be learned on a port.	All platforms	64 (static MAC locking stations) 600 (first arrival MAC locking stations)

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p>Maximum mirroring instances</p> <p>NOTE: Only two or four mirroring instance will be active at a time depending on the mirroring filter added to it. There are four hardware resource slots. Each single instance uses one such slot, while each ingress plus egress instance uses two slots. So this allows the you to use a total of four slots, while there are no more then two egress instances. The maximum possible combination for mirroring instances:</p> <ul style="list-style-type: none"> 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 <p>NOTE: The Summit X430 can only support one egress mirroring instance.</p>	All platforms	16 (including default mirroring instance)
<p>Mirroring (filters)—maximum number of mirroring filters.</p> <p>NOTE: This is the number of filters across all the active mirroring instances.</p>	BlackDiamond 8000 series BlackDiamond X8 series All Summit series E4G cell site routers	128 128 128 128
<p>Mirroring, one-to-many (filters)—maximum number of one-to-many mirroring filters.</p> <p>NOTE: This is the no. of filters across all the active mirroring instances</p>	BlackDiamond 8000 series BlackDiamond X8 series All Summit series E4G cell site routers	128 128 128 128
<p>Mirroring, one-to-many (monitor port)—maximum number of one-to-many monitor ports.</p>	All platforms	16
<p>MLAG ports—maximum number of MLAG ports allowed.</p>	BlackDiamond 8000 series BlackDiamond X8 series All Summit series, except X430 E4G cell site routers	768 768 768 768
<p>MLAG peers—maximum number of MLAG peers allowed.</p>	All platforms, except Summit X430	2

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
MPLS LDP ingress LSPs —maximum number of MPLS LSPs that can originate from a switch.	BlackDiamond 8900 xl-series	4,000
	BlackDiamond 8900-40G6X-xm	2,048
	BlackDiamond X8 series	2,048
	E4G-200	2,048
	E4G-400	4,000
	Summit X460, X480	4,000
	Summit X670, X670V-48t, X480 (40G VIM), X770	2,048
	Summit X670-G2 Summit X460-G2	2,048 4,000
MPLS LDP-enabled interfaces —maximum number of MPLS LDP configured interfaces per switch.	Summit X460, X670	32
	Summit X480, X670V-48t, X770	64
	Summit X670-G2, X460-G2	
	BlackDiamond 8900 xl-series	128
	BlackDiamond 8900-40G6x-xm	64
	BlackDiamond X8 series	64
	E4G-200, E4G-200	64
		32
MPLS LDP Sessions —maximum number of MPLS LDP sessions.	BlackDiamond 8900 xl-series	64
	BlackDiamond 8900-40G6x-xm	64
	BlackDiamond X8 series	64
	Summit X770, X670v-48t, X480	64
	Summit X670-G2, X460-G2	128
	Summit X670, X460	32
	E4G-200, E4G-400	32
MPLS LDP transit LSPs —maximum number of MPLS transit LSPs per switch.	BlackDiamond 8900 xl-series	4,000
	BlackDiamond 8900-40G6X-xm	3,000
	BlackDiamond X8 series	4,000
	E4G-200	2,700
	E4G-400	4,000
	Summit X460, X480, X770, X670V-48t, X670-G2, X460-G2	4,000
	Summit X670, X480 (VIM3-40G4x)	3,000

Table 2: Supported Limits (Continued)

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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
MSDP active peers —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
MSDP SA cache entries —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 X670-G2, X460-G2, X770	14,000
MSDP maximum mesh groups —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
Multicast listener discovery (MLD) IPv6 multicast data sender —maximum number of IPv6 multicast streams supported on a switch ^{f i} NOTE: Assumes source-group-vlan mode. For additional limits, see: <ul style="list-style-type: none"> Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode. on page 79 Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode. on page 79 	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, X460-G2	3,000
	Summit X480	3,000
	Summit X670	1,500
	Summit X770, X670-G2	14,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Multicast listener discovery (MLD) snooping per-VLAN filters —maximum number of VLANs supported in per-VLAN MLD snooping mode.	BlackDiamond e-series	250
	BlackDiamond 8800 c-series	1,000
	BlackDiamond 8900 c-series	500
	BlackDiamond 8900 xl-series	2,000
	BlackDiamond 8900-40G6X-xm	500
	BlackDiamond X8 a-series modules	500
	BlackDiamond X8 100G4X and BDx X8-100G4X-XL modules	2,000
	E4G-400, Summit X460, X460-G2	1,000
	Summit X480	2,000
	Summit X440	250
	Summit X670, E4G-200	500
Summit X770, X670-G2	1,200	
Multicast listener discovery (MLD)v1 subscribers —maximum number of MLDv1 subscribers per port ^s	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	
Multicast listener discovery (MLD)v1 subscribers —maximum number of MLDv1 subscribers per switch ^s	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	
Multicast listener discovery (MLD)v2 subscribers —maximum number of MLDv2 subscribers per port ^s	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	
Multicast listener discovery (MLD)v2 subscribers —maximum number of MLDv2 subscribers per switch ^s	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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Multicast listener discovery (MLD)v2 maximum source per group —maximum number of source addresses per group	All platforms, except Summit X430	200
Multicast VLAN registration (MVR) —maximum number of MVR senders per switch (IP multicast compression disabled). NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 series 8900-10G24X-c modules 8900-G96T-c modules 8900 xl-series 8900-40G6X-xm BlackDiamond X8 a-series modules BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules E4G-200 E4G-400 Summit X440 Summit X460, X460-G2, X480 Summit X670 VIM4-40G4x Summit X770, X670-G2	2,048 ^h 500 ⁱ 2,048 ^h 4,096 ^h 4,096 ^h 3,000 ⁱ 4,096 4,096 2,048 500 ⁱ 1,024 2,048 3,000 ⁱ 4,096
Multicast VLAN registration (MVR) —maximum number of MVR senders per switch (IP multicast compression enabled). NOTE: Assumes source-group-vlan mode. For additional limits, see: Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode. ⁱ on page 80	BlackDiamond 8800 c-series BlackDiamond 8000 e-series BlackDiamond 8900 c-series BlackDiamond 8900 xl-series BlackDiamond X8 a-series modules BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules 8900-40G6X-xm module Summit X440 Summit X460, X460-G2, E4G-400 Summit X480 Summit X670 VIM4-40G4x Summit X770, X670-G2	6,000 ⁱ 500 ⁱ 6,000 ⁱ 12,000 ^b 6,000 ⁱ 59,000 3,000 ⁱ 192 ⁱ 6,000 ⁱ 12,000 ^b 3,000 ⁱ 66,500
Network login —maximum number of clients being authenticated on MAC-based VLAN enabled ports.	BlackDiamond 8000 series (clients per module/per system) BlackDiamond X8 series Summit series	1,024 1,024 1,024

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Network login —maximum number of dynamic VLANs.	All platforms	2,000
Network login VLAN VSAs —maximum number of VLANs a client can be authenticated on at any given time.	All platforms	10
OSPFv2/v3 ECMP —maximum number of equal cost multipath OSPFv2 and OSPFv3.	All platforms, except Summit X440, X430, and E4G-200)	16
	E4G-200	8
	BlackDiamond 8800 G48Te2 (for IPv6)	N/A
OSPFv2 areas —as an ABR, how many OSPF areas are supported within the same switch.	All platforms (except X430, X440)	8
OSPFv2 external routes —recommended maximum number of external routes contained in an OSPF LSDB.	BlackDiamond 8000 series	20,000
	BlackDiamond 8900 xl-series	130,000
	BlackDiamond X8 series	20,000
	BlackDiamond BDX X8-100G4X-XL modules	130,000
	Summit X460, X670, X770, X670-G2, X460-G2	5,000
	Summit X480	130,000
	E4G-400	5,000
E4G-200	5,000	
OSPFv2 inter- or intra-area routes —recommended maximum number of inter- or intra-area routes contained in an OSPF LSDB with one ABR in OSPF domain.	BlackDiamond 8000 series	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
OSPFv2 interfaces —recommended maximum number of OSPF interfaces on a switch.	NOTE: 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
OSPFv2 links —maximum number of links in the router LSA.	All platforms, except Summit X770 and X430	400
	Summit X770	419

Table 2: Supported Limits (Continued)

Metric	Product	Limit
OSPFv2 neighbors —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
OSPFv2 routers in a single area —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
OSPFv2 virtual links —maximum number of supported OSPF virtual links.	All platforms (except X430 and X440) with Core license or higher	32
OSPFv3 areas —as an ABR, the maximum number of supported OSPFv3 areas.	All platforms (except X430 and X440) with Core license or higher	16
OSPFv3 external routes —recommended maximum number of external routes.	BlackDiamond 8000 series	10,000
	BlackDiamond X8 series	10,000
	BlackDiamond BDX X8-100G4X-XL modules	60,000
	BlackDiamond 8900 xl-series	60,000
	Summit X460, X670, X770, X670-G2, X460-G2	10,000
	Summit X480	60,000
OSPFv3 inter- or intra-area routes —recommended maximum number of inter- or intra-area routes.	E4G-400	10,000
	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	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
OSPFv3 interfaces —maximum number of OSPFv3 interfaces.	NOTE: Active interfaces only, with Advanced Edge license. (See below for Core license limits.) All platforms (except X430)	4
	NOTE: 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
OSPFv3 neighbors —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
OSPFv3 virtual links —maximum number of OSPFv3 virtual links supported.	All platforms (except X430 and X440) with Core license or higher	16
PIM IPv4 snooping —maximum number of (S,G) entries programmed in the hardware (IP multicast compression disabled). NOTE: Assumes source-group-vlan mode.	BlackDiamond 8800 c-series	2,048 ^h
	BlackDiamond 8000 e-series	500 ^h
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^h
	8900-G96T-c modules	4,096 ^h
	8900 xl-series	4,096 ^h
	8900-40G6X-xm	3,000 ⁱ
	BlackDiamond X8 a-series modules	4,096
	BlackDiamond X8-100G4X andBDX X8-100G4X-XL modules	4,096
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64
	Summit X460	2,048
Summit X480, X460-G2, X770, X670-G2	4,096	
Summit X670	3,000 ⁱ	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p>PIM IPv4 snooping—maximum number of (S,G) entries programmed in the hardware (IP multicast compression enabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p> <p>For additional limits, see:</p> <ul style="list-style-type: none"> • Layer-2 IPMC forwarding caches—(IGMP/MLD/PIM snooping) in mac-vlan mode. on page 79 • Layer-2 IPMC forwarding caches— (IGMP/MLD/PIM snooping) in mixed-mode. on page 79 	BlackDiamond 8800 c-series	6,000 ⁱ
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900 c-series	6,000 ⁱ
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ⁱ
	BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules	59,000 ⁱ
	E4G-200	3,000 ⁱ
	E4G-400	6,000 ⁱ
	8900-40G6X-xm	3,000 ⁱ
	Summit X440	192 ⁱ
	Summit X480	12,000 ^b
	Summit X460	6,000 ⁱ
	Summit X670	3,000 ⁱ
Summit X770, X670-G2	66,500	
Summit X460-G2	21,000	
<p>PIM IPv4—maximum routes—maximum number of (S,G) entries installed in the hardware (IP multicast compression disabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p>	BlackDiamond 8800 c-series	2,048 ^h
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^h
	8900-G96T-c modules	4,096 ^h
	8900 xl-series	4,096 ^h
	8900-40G6X-xm	3,000 ⁱ
	BlackDiamond X8 a-series modules	4,094
	BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules	4,096
	E4G-200	2,048
	E4G-400	2,048
	Summit X440	64 ⁱ
	Summit X480, X670-G2, X460-G2	4,096
Summit X460	2,048	
Summit X670	3,000 ⁱ	
Summit X770	4,096	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p>PIM IPv4—maximum routes—maximum number of (S,G) entries installed in the hardware (IP multicast compression enabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p> <p>For additional limits, see:</p> <p>Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode.ⁱ on page 80</p>	BlackDiamond 8800 c-series	6,000 ⁱ
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900 c-series	6,000 ⁱ
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ^f
	BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules	59,000 ^f
	E4G-200	3,000 ⁱ
	E4G-400	6,000 ⁱ
	8900-40G6X-xm modules	3,000 ⁱ
	Summit X440	192
	Summit X480	12,000 ^b
	Summit X460	6,000 ⁱ
	Summit X670	3,000 ⁱ
Summit X770, X670-G2	66,500	
Summit X460-G2	21,000	
<p>PIM IPv4-SSM (maximum SSM routes)—maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multicast compression disabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p>	BlackDiamond 8800 c-series	2,048 ^h
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900 series	
	8900-10G24X-c modules	2,048 ^h
	8900-G96T-c modules	4,096 ^h
	8900 xl-series	15,000
	8900-40G6X-xm	3,000 ⁱ
	BlackDiamond X8 a-series modules	4,094
	BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules	4,094
	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 ⁱ	
Summit X770	4,096	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
<p>PIM IPv4-SSM (maximum SSM routes)—maximum number of (S,G) entries installed in the hardware with PIM SSM configuration (IP multicast compression enabled).</p> <p>NOTE: Assumes source-group-vlan mode.</p> <p>For additional limits, see:</p> <p>Layer-3 IPMC forwarding caches—(PIM, MVR, PVLAN) in mixed-mode.ⁱ on page 80</p>	BlackDiamond 8800 c-series	6,000 ⁱ
	BlackDiamond 8000 e-series	500 ⁱ
	BlackDiamond 8900 c-series	6,000 ⁱ
	BlackDiamond 8900 xl-series	12,000 ^b
	BlackDiamond X8 a-series modules	6,000 ⁱ
	BlackDiamond X8-100G4X and BDX X8-100G4X-XL modules	59,000 ⁱ
	E4G-200	3,000 ⁱ
	E4G-400	6,000 ⁱ
	8900-40G6X-xm	3,000 ⁱ
	Summit X440	192 ⁱ
	Summit X480	12,000 ^b
	Summit X460	6,000 ⁱ
	Summit X670	3,000 ⁱ
Summit X770, X670-G2	66,500	
Summit X460-G2	21,000	
<p>PIM IPv6 (maximum routes)—maximum number of (S,G) entries installed in the hardware.</p> <p>NOTE: Assumes source-group-vlan mode.</p>	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 and BDX X8-100G4X-XL modules	30,000 ^e
	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	30,000	
<p>PIM IPv4 (maximum interfaces)—maximum number of PIM active interfaces.</p>	All platforms, except Summit X430 and X440	512
	Summit X440	253
<p>PIM IPv4 (maximum interfaces)—maximum number of PIM snooping enabled interfaces.</p>	All platforms, except Summit X430	512
<p>PIM IPv4 Limits—maximum number of multicast groups per rendezvous point</p>	All platforms, except Summit X430	180

Table 2: Supported Limits (Continued)

Metric	Product	Limit
PIM IPv4 Limits —maximum number of multicast sources per group	All platforms, except Summit X430	175
PIM IPv4 Limits —maximum number of dynamic rendezvous points per multicast group	All platforms, except Summit X430	145
PIM IPv4 Limits —static rendezvous points	All platforms, except Summit X430	32
PIM IPv6 (maximum interfaces) —maximum number of PIM active interfaces	All platforms, except Summit X430	512
PIM IPv6 Limits —maximum number of multicast group per rendezvous point	All platforms, except Summit X430	70
PIM IPv6 Limits —maximum number of multicast sources per group	All platforms, except Summit X430	43
PIM IPv6 Limits —maximum number of dynamic rendezvous points per multicast group	All platforms, except Summit X430	64
PIM IPv6 Limits —maximum number of secondary address per interface	All platforms, except Summit X430	70
PIM IPv6 Limits —static rendezvous points	All platforms, except the Summit X430	32
Policy-based routing (PBR) redundancy —maximum number of flow-redirects.	All platforms	256 ^t
Policy-based routing (PBR) redundancy —maximum number of next hops per each flow-direct.	All platforms	32 ^t
Port-specific VLAN tags —maximum number of port-specific VLAN tags	All platforms	1,023

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Port-specific VLAN tags —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	
Private VLANs —maximum number of subscribers. Assumes a minimum of one port per network and subscriber VLAN.	BlackDiamond 8800 c-, e-, xl-series with eight modules of 48 ports	383
	8900-G96T-c modules	767
	BlackDiamond X8 series	103
	Summit X770	63
	Summit X670-G2, X670v-48t	47
	Summit X670	23
	Summit X480	53
	Summit X460-G2, X460	25
	Summit X440	27
	Summit X430	11
	E4G-200	33
E4G-400		
Private VLANs —maximum number of private VLANs with an IP address on the network VLAN. NOTE: This limit is dependent on the maximum number of private VLANs in an L2-only environment if the configuration has tagged and translated ports.	Summit X770, X670-G2, X460-G2	1,024
	Summit X670, X480, X460, X460, X480	512
	Summit X440	127
	E4G-200, E4G-400	512

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Private VLANs —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	597
	Summit X670	597
	Summit X460	820
	Summit X770, X670-G2, X460-G2	1,280
PTP/1588v2 Clock Ports	Summit X770, X460-G2, X670-G2, and E4G-200, E4G-400 cell site routers	32 for boundary clock
		1 for ordinary clock
PTP/1588v2 Clock Instances	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	2 combinations: <ul style="list-style-type: none"> Transparent clock + ordinary clock Transparent clock + boundary clock
PTP/1588v2 Unicast Static Slaves	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	40 entries per clock port
PTP/1588v2 Unicast Static Masters	Summit X770, X670-G2, X460-G2, and E4G-200, E4G-400 cell site routers	10 entries per clock type
Route policies —suggested maximum number of lines in a route policy file.	All platforms	10,000
RIP Learned Routes —maximum number of RIP routes supported without aggregation.	All platforms, except Summit X430	10,000
RIP neighbors —maximum number of RIP neighbors.	E4G-200	256

Table 2: Supported Limits (Continued)

Metric	Product	Limit
RIP interfaces on a single router —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	
RIPng learned routes —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	
Spanning Tree (maximum STPDs) —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
Spanning Tree PVST+ —maximum number of port mode PVST domains. NOTE: <ul style="list-style-type: none"> Maximum of 10 active ports per PVST domain when 256 PVST domains are configured. Maximum of 7 active ports per PVST domain when 128 PVST domains are configured. 	BlackDiamond X8 and 8900 series switches	256
	Summit X670, X770, X670-G2	256
	Summit X460, X480, X440, X460-G2	128
	Summit X430	50
	E4G-400	128
Spanning Tree —maximum number of multiple spanning tree instances (MSTI) domains.	All platforms (except Summit X430 and Summit X440)	64
	Summit X440	32
	Summit X430	5

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Spanning Tree —maximum number of VLANs per MSTI. NOTE: Maximum number of 10 active ports per VLAN when all 500 VLANs are in one MSTI.	BlackDiamond X8	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	
Spanning Tree —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	
Spanning Tree (802.1d domains) —maximum number of 802.1d domains per port.	All platforms	1
Spanning Tree (number of ports) —maximum number of ports including all Spanning Tree domains.	All platforms (except Summit X430 and Summit X440)	4,096
	Summit X440	2,048
	Summit X430	1,024
Spanning Tree (maximum VLANs) —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	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
SSH (number of sessions) —maximum number of simultaneous SSH sessions.	All platforms	8
Static MAC multicast FDB entries —maximum number of permanent multicast MAC entries configured into the FDB.	BlackDiamond 8000 c-, e-, xl-series	1,024
	BlackDiamond X8 series	1,024
	All Summits E4G-200, E4G-400	1,024
Syslog servers —maximum number of simultaneous syslog servers that are supported.	All platforms	4
Telnet (number of sessions) —maximum number of simultaneous Telnet sessions.	All platforms	8
TRILL —trees rooted from switch	BlackDiamond X8	1
	Summit X670, X770	1
TRILL —computed trees	BlackDiamond X8	1
	Summit X670, X770	1
TRILL —TRILL VLANs	BlackDiamond X8	4
	Summit X670, X770	4
TRILL —forwarding VLANs	BlackDiamond X8	4,095
	Summit X670, X770	4,095
TRILL —forwarding ports	BlackDiamond X8	All
	Summit X670, X770	All
TRILL —RBridge FDB entries	BlackDiamond X8	128,000
	Summit X670	128,000
	Summit X770	288,000
TRILL —ECMP RBridge next hops	BlackDiamond X8	8
	Summit X670, X770	8
TRILL —neighbor adjacencies	BlackDiamond X8	32
	Summit X670, X770	32
TRILL —nodes	BlackDiamond X8	256
	Summit X670, X770	256
TRILL —links	BlackDiamond X8	2,000
	Summit X670, X770	2,000

Table 2: Supported Limits (Continued)

Metric	Product	Limit
Virtual routers —maximum number of user-created virtual routers that can be created on a switch. NOTE: Virtual routers are not supported on Summit X440 series switches.	BlackDiamond 8000 c-, xl-, xm-series	63
	BlackDiamond X8 series	63
	E4G-200, E4G-400	63
	Summit X460, X460-G2, X480, X670, X670-G2, X770	63
Virtual router forwarding (VRFs) —maximum number of VRFs that can be created on a switch. NOTE: Subject to other system limitations.	All platforms, except Summit X440, X430	960 *
Virtual router protocols per VR —maximum number of routing protocols per VR.	All platforms, except Summit X440, X430	8
Virtual router protocols per switch —maximum number of VR protocols per switch.	All platforms, except Summit X440, X430	64
VLAN aggregation —maximum number of port-VLAN combinations on any one superVLAN and all of its subVLANs.	All platforms (except Summit X430, X440)	1,000
	Summit X440, X430	256
VLANs —includes all VLANs. NOTE: ExtremeXOS supports only 4,092 user-configurable VLANs. (VLAN 1 is the default VLAN, and 4,095 is the management VLAN, and you may not configure them.)	All platforms	4,094
VLANs —maximum number of port-specific tag VLANs.	BlackDiamond 8800 xl-series only, BlackDiamond X8 series	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
VLANs —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

Table 2: Supported Limits (Continued)

Metric	Product	Limit
VLANs (Layer 2) —maximum number of Layer 2 VLANs.	All platforms	4,094
VLANs (Layer 3) —maximum number of VLANs performing IPv4 and/or IPv6 routing. Excludes sub-VLANs.	Summit X460-G2, X670, X770, X670-G2, and BlackDiamond X8	2,048
	Summit X440	254
	Summit X480, X460	512
	E4G-200, E4G-400	512
VLANs (maximum active port-based) —(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	1
VLANs (maximum active protocol-sensitive filters) —number of simultaneously active protocol filters in the switch.	All platforms	15
VLAN translation —maximum number of translation VLANs. Assumes a minimum of one port per translation and member VLAN.	BlackDiamond 8000 a-, c-, e-, xl series	with eight modules of 48 ports (383) 8900-G96T-c modules (767)
	Summit X770	103
	Summit X670-G2, X670v-48t	63
	Summit X670	47
	Summit X480	53
	Summit X460-G2	53
	Summit X460	57
	E4G-200	11
	E4G-400	33
	Summit X440	25
Summit X430	27	

Table 2: Supported Limits (Continued)

Metric	Product	Limit
VLAN translation —maximum number of translation VLAN pairs with an IP address on the translation VLAN. NOTE: This limit is dependent on the maximum number of translation VLAN pairs in an L2-only environment if the configuration has tagged and translated ports.	Summit X770, X670-G2	1,024
	Summit X670, X480, X460	512
	Summit X460-G2	1,024
	E4G-200, E4G-400	512
	Summit X440	127
VLAN translation —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	2,000
	E4G-400, E4G-200	2,000
	Summit X440, X430	512
	Summit X480, X670, X770, X670-G2, X460-G2	2,046
VRRP (v2/v3-IPv4) (maximum instances) —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
VRRP (v3-IPv6) (maximum instances) —maximum number of VRRP instances for a single switch, with Advanced Edge license or higher. (VRRP-VRRPv3-IPv6)	BlackDiamond X8, 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	255
	Summit X440	15
VRRP (v2/v3-IPv4/IPv6) (maximum VRID) —maximum number of unique VRID numbers per switch.	All platforms with Advanced Edge license or higher, except Summit X430	7
VRRP (v2/v3-IPv4/IPv6) (maximum VRIDs per VLAN) —maximum number of VRIDs per VLAN.	All platforms with Advanced Edge license or higher, except for Summit X430	7

Table 2: Supported Limits (Continued)

Metric	Product	Limit
VRRP (v2/v3-IPv4/IPv6) (maximum ping tracks) —maximum number of ping tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
VRRP (maximum ping tracks) —maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	All platforms, except the Summit X440 Summit X440 Hello interval: 20 centiseconds Hello interval: 1 second	8 (20 centisecond or 1 second hello interval) 1 1
VRRP (v3-IPv6) (maximum ping tracks) —maximum number of ping tracks per VRRP Instance under 128 VRRP instances, with Advanced Edge license or higher.	All platforms, except the Summit X440 Summit X440 Hello interval: 20 centiseconds Hello interval: 1 second	8 (20 centisecond or 1 second hello interval) 1 (IPv6) 1 (IPv6)
VRRP (v2/v3-IPv4/IPv6) (maximum iproute tracks) —maximum number of IP route tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
VRRP (v2/v3-IPv4/IPv6) —maximum number of VLAN tracks per VLAN.	All platforms with Advanced Edge license or higher, except Summit X430	8
XML requests —maximum number of XML requests per second. NOTE: Limits are dependent on load and type of XML request. These values are dynamic ACL data requests.	BlackDiamond 8800 c-series with 100 DACLs with 500 DACLs BlackDiamond 8900 series with 100 DACLs with 500 DACLs Summit X480, X670 with 100 DACLs with 500 DACLs	10 3 10 3 4 1
XNV authentication —maximum number of VMs that can be processed (combination of local and network VMs).	All platforms, except Summit X430	2,048
XNV database entries —maximum number of VM database entries (combination of local and network VMs).	All platforms, except Summit X430	16,000
XNV database entries —maximum number of VPP database entries (combination of local and network VPPs).	All platforms, except Summit X430	2,048

Table 2: Supported Limits (Continued)

Metric	Product	Limit
XNV dynamic VLAN —Maximum number of dynamic VLANs created (from VPPs /local VMs).	All Platforms, except Summit X430	2,048
XNV local VPPs —maximum number of XNV local VPPs.	All platforms, except Summit X430 Ingress Egress	2,048 512
XNV policies/dynamic ACLs —maximum number of policies/dynamic ACLs that can be configured per VPP. ^u	All platforms, except Summit X430 Ingress Egress	8 4
XNV network VPPs —maximum number of XNV network VPPs. ^u	All platforms, except Summit X430 Ingress Egress	2,048 512

- a. The table shows the total available.
- b. Limit depends on setting configured for `configure forwarding external-tables`.
- c. When there are BFD sessions with minimal timer, sessions with default timer should not be used.
- d. Based on in “none more-I2” mode.
- e. Based on forwarding internal table configuration “more I2”.
- f. Effective capacity varies based on actual MAC addresses and VLAN IDs used and hash algorithm selected.
- g. Based on “I2-only mode”.
- h. Applies only if all enabled BlackDiamond 8000 I/O modules are BlackDiamond 8000 c-, xl-, or xm-series modules.
- i. 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.
- j. 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.
- k. Based on forwarding internal table configuration “more I3-and-ipmc”.
- l. Based on forwarding external table configuration “none” and forwarding internal table configuration “more I3-and-ipmc”.
- m. Based on forwarding external table configuration “I3-only ipv4”.
- n. The limit depends on setting configured with `configure iproute reserved-entries`.
- o. Based on forwarding external table configuration “I3-only ipv4”.
- p. Based on forwarding external table configuration “I3-only ipv6”.
- q. Based on forwarding internal table configuration “I2-and-I3”.
- r. The IPv4 and IPv6 multicast entries share the same hardware tables, so the effective number of IPv6 multicast entries depends on the number of IPv4 multicast entries present and vice-versa.
- s. If IGMP and MLD are simultaneously configured on the switch, the number of effective subscribers supported would be appropriately lessened.
- t. Sum total of all PBR next hops on all flow redirects should not exceed 1024.
- u. 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 106](#)
- [Known Behaviors on page 110](#)
- [Resolved Issues in ExtremeXOS 15.7.3-Patch1-6 on page 111](#)
- [Resolved Issues in ExtremeXOS 15.7.3-Patch1-3 on page 113](#)
- [Resolved Issues in ExtremeXOS 15.7.3 on page 117](#)
- [Resolved Issues in ExtremeXOS 15.7.2 on page 121](#)
- [Resolved Issues in ExtremeXOS 15.7.1 on page 129](#)

Open Issues

The following are new open issues for supported features found in ExtremeXOS 15.7.3-Patch1-6.

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

CR Number	Description
General	
xos0057231	An FDB entry created by ARP with "i" flag set is not removed from the FDB table after a static entry for the same IP address is added with a different MAC value.
xos0053450	Bootp_Relay with Option 82: Client connected VLAN with more than one port with option on and check "on" and policy "keep". Check is not working.
xos0057297	When MD5 password is different for both broadcast server and client, the NTP session is established anyway, and no warning is indicated in "leap indicator".
xos0057356	The output of the <code>show ntp association 1.1.1.1 statistics</code> command shows the local interface IP address as NTP association IP address or 0.0.0.0. Workaround: Running <code>clear counter</code> command corrects the IP address.
xos0057392	NTP session still shows synchronized after disabling port in NTP client interface.
xos0057517	With global NTP server configured and an NTP session established between server and broadcast client, after disabling NTP, the peer loses its active peer status (*), but for the broadcast client, status (*) is not removed.

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

CR Number	Description
BlackDiamond X8 Series Switches	
xos0059214	GRE tunnel traffic for some destinations is slow-path forwarded after running two failovers.
xos0059633	On L2VPN session between BlackDiamond X8 series switches with traffic established over RSVP primary path, restarting MPLS on Provider Edge (PE) switches produces the following MPLS kernel errors: 01/08/2015 08:37:13.64 <Error:Kern.MPLS.Error> Slot-5: extreme_mpls_del_tid_vp_nh_xref: MPLS TID+VP->NH tid = 1, vp : 0x4 map entry not found in btree rv = -7 (Entry not found)
xos0059927	For the 100G4X and 100G4X-XL modules, tagged data traffic is not going to untagged service VMAN after deleting tagged service VMAN from the VPLS instance, and then creating new untagged service VMAN and attaching it to the VPLS.
BlackDiamond 8800 Series Switches	
xos0057741	Deleting a shared group from BlackDiamond 8800 series switches deletes the port linked to MVRP, but the dynamic VLAN still holds the port.
Summit Family Switches	
xos0058437	For Summit X460 and X670-G2 series switches, the buffer for Weighted Random Early Detection (WRED) queues is incorrectly allocated at 10% of shared memory plus minimum guarantee, when it should be 100% of shared memory plus minimum guarantee.
xos0058868	Configuring SRP for a port after disabling VLAN containing the port, puts the master port of SRP in unblocking state. FDB is learned through that port, though VLAN is disabled. After un-configuring SRP for the ports present in the VLAN, this unblocks redundant port also, causing loop in the network.
xos0059373	Gigabit Ethernet compliance of 100LX10 optics appears as "UNKNOWN" in <code>debug hal show optic-info</code> command.
xos0057605	Summit X460 and X460-G2 series switches and E4G-200 cell site routers have more than one option slot. For these platforms, entPhysicalDesc Object in entPhysicalTable returns "Option Slot-1" for two option-card slots.
xos0059345	When 100LX10,100FX without PHY optics are inserted in 10G SFP+ ports, no warning log messages occur to indicate that these optics are not supported on 10G SFP+ ports.

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

CR Number	Description
SummitStack	
xos0060218	In Summit X670-G2 stacks, with bi-directional L2 multicast traffic running, running failover produces the following error message: <Erro:HAL.LAG.CfgFail> Slot-2: Failed to configure link aggregation group 1:10 on slot 1 unit 0: Operation timed out
xos0058407	On SummitStacks with Summit X670 and X670 as MLAG peers and connected over LACP LAG, and L2VPN is configured with primary Explicit Route Option (ERO) and secondary ERO. While traffic is flowing, performing a slot failover produces error messages.
xos0058637	On Summit X670 stacks, the following MPLS kernel errors occur when after issuing the command <code>delete VPLS all</code> : "Slot-1 BA-CO-NSNP_75.17 # Oct 24 02:03:23 BA-CO-NSNP_75 <Erro:Kern.MPLS.Error> (ems_main.c:446) bcm_custom_extr_vfp_action_mpls_inport_set(remove) failed for unit = 0 vpn = 0x65, port = 0x800000f vp = 0x18000191, rv = -13 (Invalid identifier)"
Summit X460-G2 Series Switches	
xos0059395	Summit X460-G2-1G switches do not produce warning log messages (not compatible) when 100LX10 optics are inserted in the highest numbered 1G ports.
xos0055999	For Summit X670-G2 and X460-G2 series switches, currently the ENTITY-MIB:entPhysicalsFRU for fan and power modules shows the values as FALSE. However since these fan modules and power modules are replaceable units, this should be shown as TRUE.
BGP	
xos0058441	After creating a BGP peering session between link local IPv6 addresses with the scope ID specified, deleting the VLAN containing link local IPv6 address. and then issuing the command <code>show configuration bgp</code> , switch reboots with "Epm application wdg timer warning" error message.
xos0058666	BGP multicast routes are not displayed in routing table when origin is specified.
MPLS	
xos0057448	After adding more than two or three Explicit Route Option (ERO) Excludes to primary path, the command <code>show mpls rsvp-te lsp my_lsp1 detail</code> shows "Path Computation Failure" for the primary path and is using the secondary path as observed in the command <code>traceroute mpls lsp my_lsp1</code> .
xos0057509	When primary path is configured with strict source and added link with exclude option, label-switched path (LSP) comes up. In secondary path adding Explicit Route Option (ERO) with loose source and excluding the same link, LSP never comes up.
xos0057834	After restarting MPLS, disabling primary path makes secondary path go down.

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

CR Number	Description
xos0058421	VPWS (tag include): Outer dot1q tag is missing from frame received on remote Provider Edge (PE) switch.
xos0057510	Traffic is dropped on switches after disabling primary when RSVP label-switched path (LSP) is fast reroute enabled. Workaround: Issue command <code>restart process mpls</code> .
xos0058974	After deleting, and then adding VPWS tunnel (rapidly, by using script) the <code>mpls statistics l2vpn</code> counters stops incrementing.
xos0059530	Disabling, and then enabling MPLS on primary path P node is causing data traffic loss at ingress of a VPLS Provider Edge (PE) switch for 15 seconds. Sometimes this causes VPLS to go into signal state, which can cause traffic loss. P node should tear down the LSP/primary path before disabling MPLS/RSVP, so that LSP can move to secondary path quickly.
RIP	
xos0058173	Route policy created in a RIP VLAN to permit some routes in network range 192.168.0.0/24, denies all the RIP routes learned via that particular VLAN interface.

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 Requests (CRs)

CR Number	Description
General	
xos0058855	MLAG-PIM-SSM: Traffic loss for a subscriber occurs when another subscriber leaves on a two-peer MLAG setup. Workaround: Increase the leave-timeout and last membership query interval to 2 seconds.
xos0055791	Block size option is not available in bootROM download image command. Generally an image is downloaded from bootROM either when the switch is newly purchased or for an emergency recovery (for example, when switch is constantly rebooting without getting a command line prompt), so only minimal arguments are provided for the command (for example, IP address, gateway, and download image location).
xos0057718	SCP2 file transfer from external SCP client does not work. OpenSSH uses the RCP protocol, which is disabled in the ExtremeXOS for security reasons. Consequently, OpenSSH SCP does not work with the ExtremeXOS SSH implementation. Workaround: To transfer files, use SFTP instead of SCP from external clients.
xos0058944	PVLAN network-tagged packets are allowed to ingress on subscriber VLAN ports.
BlackDiamond 8800 Series Switches	
xos0058703	IPv6 ECMP is not supported on BlackDiamond 8800 G48Te2 modules.
Summit Family Switches	
xos0059482	Configuring QoSprofile X to "0" using peak_rate or maxBw does not work. All traffic is forwarded as if QoSprofile X is set to 100% maxBw.
IP Routing Protocols	
xos0057722	Duplicated traffic is seen when SPT threshold is set to 'infinity' on WMLAG peers Workaround: Do not set SPT threshold at 'infinity' in MLAG and WMLAG topologies. Since no (S;G)s are created, no checkpoint occurs and this results in duplicated traffic.

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

CR Number	Description
MPLS	
xos0058780	On Summit X460-G2 and X670-G2 series switches, and BlackDiamond X8 switches with XB-100G4X-XL modules, the command <code>show mpls statistics l2vpn</code> displays double the TxPackets count as the actual packets transmitted over the VPLS.
OSPF	
xos0057100	The command <code>unconfigure ospfv3</code> is not removing the VLANs from the OSPFv3 domain. All the neighbors are still visible.
STP	
xos0058362	Ports configured as auto with auto-edge feature turned on, do not have this status correctly shown in <code>show <stpd> port</code> command. Port operation mode appears as "Point-point". Workaround: Port operation mode appears correctly in <code>show <stpd> port detail</code> command.

Resolved Issues in ExtremeXOS 15.7.3-Patch1-6

The following issues were resolved in ExtremeXOS 15.7.3-Patch1-6. ExtremeXOS 15.7.3-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 12.6.3, 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-patch1-9, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.5.2, ExtremeXOS 15.6.4 and ExtremeXOS 15.7.3 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
General	
xos0062850	When upgrading ExtremeXOS to 15.7 or later releases, the web HTTP access is enabled even though it is disabled in the configuration.
xos0063815	Traffic is dropped on one VPLS instance when another VPLS instance running on same switch is deleted.
xos0061462	Traffic forwarding on specific ports in a VLAN is affected when another port in same VLAN is configured for port specific tag.
xos0057785	STP domain tag is removed when all ports are deleted from STP auto-bind enabled-VLANs.

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

CR Number	Description
xos0063134	Traffic stops after disabling, and then enabling LAG portst having pstag with static FDB.
xos0063240	ACL process ends unexpectedly when switch has clear-flow ACL rule with count interval greater than snmptrap generation timer.
xos0063245	With IGMP per-VLAN mode, VRRP flaps occur after adding tagged ports to VLANs.
xos0063258	OSPFv3 routes are not updated in routing tables if there is a delay in receiving LSAs.
xos0063271	Layer 3 packets in non-default virtual routers are slow-path forwarded after disabling MPLS in the peer switch.
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.
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.

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

CR Number	Description
SummitStack	
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.
xos0062949	HAL process ends unexpectedly in stack after executing the following commands: <pre>debug hal configure stacking pdu-trace mask 0xf debug hal configure stacking pdu-trace capture cap_file</pre>
Summit X460-G2 Series Switches	
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.
BlackDiamond 8800 Series Switches	
xos0063333	In BlackDiamond 8800 series switches, optics information is not detected and ports remain in "Ready" state after reboot.
xos0063614	Kernel crash occurs when receiving DHCP packets with invalid field values.

Resolved Issues in ExtremeXOS 15.7.3-Patch1-3

The following issues were resolved in ExtremeXOS 15.7.3-Patch1-3. ExtremeXOS 15.7.3-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 12.6.3, 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-patch1-9, ExtremeXOS 15.4.2.8, ExtremeXOS 15.5.5.2, ExtremeXOS 15.6.4 and ExtremeXOS 15.7.3 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
General	
xos0061781	Identity manager entries become stale when clients are moved from one port to another in sub-VLANs.
xos0062367	ACL process ends unexpectedly on repeated refresh of ACL policy with clear-flow action.

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

CR Number	Description
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
xos0050590	Multicast cache entries are not programmed in hardware due to false resource full condition due to IPv6 cache entries.
xos0057269	SNMP trap extremelpSecurityViolation is sent with incorrect VLAN description.
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.
xos0059569	OSPFv3 external routes flap and traffic loss occurs when introducing new ABR which has reachability to destination.
xos0061027	For SummitStacks, creating or deleting non-default QoS profiles may cause some ports to flap.
xos0061505	After a topology change in the network, BGP routes requiring two levels of recursive lookup are programmed in hardware with incorrect next hops.
xos0061507	SNMPget on EXTREME-SOFTWARE-MONITOR table returns value with incorrect OID.
xos0061788	The process devmgr ends unexpectedly during snmpwalk when continuous EMS logs are sent to the switch console.
xos0061855	Configured OSPF neighbor is not retained after rebooting.
xos0062045	LLDP packets are tunnelled over L2VPN.
xos0062240	Port that was administratively disabled becomes up after enabling rx pause.
xos0062366	After reboot, DHCP binding entries are not restored via vr-default.
xos0062441	The process rtMgr ends unexpectedly when IPv6 static route is deleted.
xos0062537	HAL crash occurs when redirect-port-list action contains more than 64 ports.

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

CR Number	Description
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.
xos0062728	OSPFv3 best path is not selected after issuing <code>restart ports all</code> in peer switch.
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.
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.
xos0063368	In an MLAG configured switch, FDBs are not installed in hardware after reboot if there are frequent MACMoves between MLAG port and ISC.
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.
xos0057538	OSPFv3 fails to select the best cost external route.
xos0058808	Rarely, MAC addresses of authenticated clients learned on NetLogin-enabled ports are not programmed in hardware.
xos0059341	OSPFv3 external routes are flushed after interface cost is changed if multiple ABRs are present for an area.
xos0059446	OSPFv3 external routes are flushed when ports from OSPFv3 user VR are deleted.
xos0060463	OSPFv3 external routes are flushed after the command <code>restart ports all</code> is executed in area border router.
xos0062255	CEP CVID configurations is missing after adding/deleting the port from sharing.
xos0062494	Source MAC addresses learned through MVRP packets on a blocked port (STP) cause traffic to be dropped.
xos0062504	You can set a GPTP "peer delay current interval" outside of the correct range of -3 to 17..

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

CR Number	Description
xos0062508	When a port is added in a loopback VLAN, OSPFv3 route is not advertised with /128 mask.
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.
xos0062756	Output of <code>show network-clock gtp current-set</code> command shows incorrect timestamp for "Last GM Change Event".
Summit X440 Series Switches	
xos0062621	On Summit X440-8p switches, the <code>show fan</code> command output displays that the fan is unsupported.
Summit X460 Series Switches	
xos0063206	Cannot add L2 entries in hardware due to a full L2 table caused by hash collisions.
Summit X460-G2 Series Switches	
xos0060559	Hot-swapping fan modules in Summit X460-G2 series switches puts some fans in failed state.
xos0062855	On the Summit X460-G2 series switches, VPLS packets are forwarded with two tags when the service VLAN ports are also members of an untagged VMAN.
Summit X480 Series Switches	
xos0061071	In Summit X480-48t switches, optic-info is not detected for combo ports.
Summit X670 Series Switches	
xos0059514	On Summit X670V-48x switches, after multiple reboots, 40G ports can remain in ready state, rather than coming to active state, which impacts the traffic passing through those ports.
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.
xos0063137	Known unicast traffic is not shared between the stacking high-gig trunk ports.
SummitStack	

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

CR Number	Description
xos0061791	On SummitStacks containing master and standby nodes of different switches, the standby node may go to failed state after a node reboot.
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.
xos0063349	Switch stops responding to SNMP requests if SNMP get for multiple OIDs is continuously initiated.
xos0063242	Stacks configured as DHCP clients do not respond to pinging after failover.
BlackDiamond 8800 Series Switches	
xos0062009	In BlackDiamond 8800 series switches with XL modules, clearing FDBs when there is a loop causes the FDBs to lose synchronization across slots or switching units.

Resolved Issues in ExtremeXOS 15.7.3

The following issues were resolved in ExtremeXOS 15.7.3. ExtremeXOS 15.7.3 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, ExtremeXOS 15.7.1, and ExtremeXOS 15.7.2. 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
General	
xos0056885	IGMPv3 packets with invalid checksum are not dropped.
xos0058514	SnmpEngineBoots value appears as a negative value on certain switches.
xos0059448	The process xmlc ends unexpectedly due to incorrect accessing of XML target.
xos0060622	The following error message appears after disable/enable sharing: <pre>extreme_mpls_del_vp_nh_xref: MPLS ECMP NH vp : 0xddf 13intf = 0xebf map entry not found in btree rv = -7 (Entry not found)</pre>

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

CR Number	Description
xos0060792	SNMP authentication failure log message and trap is inappropriately generated when switch detects "Not In Time Windows" error.
xos0061092	Traffic forwarding on VPLS-serviced VMAN stops after link flap.
xos0061229	FDB entries may be learned incorrectly when a port that is a member of a tagged VPLS VLAN has the final untagged VMAN removed from the same port.
xos0061495	STP auto-bind configuration does not appear in <code>show configuration</code> command after upgrade.
xos0062017	DHCP trusted port configuration is lost after disabling, and then re-enabling LAG.
xos0062128	L3VPN traffic is not forwarded after executing <code>disable port</code> and <code>enable port</code> in MPLS core network.
xos0062143	The command <code>disable netlogin mac</code> does not allow traffic on the port.
xos0062427	EDP process ends unexpectedly when CDP packets without portId TLV are received.
xos0062428	Member ports with a modified speed configuration that is different than the master port should not be allowed in LAG.
xos0062499	Multicast packets are dropped in Layer 2 bridged VLANs.
xos0062719	Allow use of 3rd-party optics without any additional license.
xos0062789	Disabling learning on LAG ports does not flush FDB entries.
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).
xos0062914	The process <code>mcmgr</code> ends unexpectedly after receiving corrupted IGMPv3 join packets on MLAG ports.
xos0059942	SSH connection ends when <code>show</code> commands produce lengthy output.
xos0061198	Disabling VPN-VRF affects traffic on another VPN-VRF.
xos0061385	EAPS process ends unexpectedly after deleting EAPS shared-port configuration.
xos0062260	BGP process ends unexpectedly when local address or password is changed for BGP neighbor, and then you immediately execute a BGP <code>show/configuration</code> command.
xos0062271	CLI memory leak occurs when executing <code>show</code> commands with <code>include</code> option through script.
xos0062300	CEP CVID ranges, other than first VLAN, fail when access port is a trunk.

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

CR Number	Description
xos0062301	Packet drops occurs between customer edge switches when VMAN and CVID tag are the same.
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.
xos0062705	Kernel oops can occur after clearing IPMC FDB
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.
xos0061886	SNMP master process ends unexpectedly with signal 6 with certain sequence of snmpbulkget and snmpget.
xos0062113	The <code>show power</code> command output does not display power usage for PSUs with part numbers starting with "800515"
Summit X430 Series Switches	
xos0059486	On Summit X430 series switches, optics are not detected after repeated removal and reinsertion of optics when CPU is busy.
xos0061864	FDB process consumes more than 20% utilization when Summit X430 switches are configured with 100+ VLANs.
BlackDiamond 8800 Series Switches	
xos0055892	Errors appear indicating that the SFP/GBIC EEPROM cannot be read for some SFPs optics: <pre>&lt;Error:Card.Error> Slot-5: aspendiags: Got bad rv (-1) reading XFP Diagnostic Monitor EEPROM for slot 5, port 8 &lt;Error:Card.Error> Slot-5: aspendiags: Unable to open XFP EEPROM for port 8, rv -1</pre>
xos0061748	Error message "aspendiags: Unable to read GBIC EEPROM" appears during bootup of c-series I/O modules.
xos0061994	Packets are not forwarded over VPWS after rebooting BlackDiamond 8800 switches.
xos0062060	FDB entries are not programmed in hardware even though hardware resources have sufficient capacity.
xos0062535	On BlackDiamond 8900-10G24Xc I/O modules, packets are dropped by ACL rules with "redirect-port-list" action.

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

CR Number	Description
BlackDiamond X8 Switches	
xos0056773	Resetting a connected I/O fabric module, produces the following error messages: 04/07/2014 10:09:10.07 <Error:Kern.Card.Error> Slot-7: pca9506_read, not enough buffer provided (need 5 bytes, have 1 bytes) 04/07/2014 10:09:10.57 <Error: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 <Error: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.
xos0061186	Bytes counter associated with <code>show port utilization</code> command output displays inaccurate value for 100G ports when utilization exceeds 40%.
xos0061639	Packets ingressing on VLAN-bridged interfaces (Layer 2 VLANs) are not forwarded when the destination MAC address is the same as the switch MAC address, and the switch has at least one Layer 3 interface.
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.
SummitStack	
xos0062570	In SummitStacks, executing the command <code>enable sflow ports</code> all enables sFlow inappropriately on stacking ports.
xos0061067	On SummitStacks with PoE and non-PoE switches, executing the command <code>show inline-power slot <slot_number></code> produces an XML error.
xos0062522	In Summit stacking switches, standby slots go to failed state when very large number of log messages are continuously generated in the switch.
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.
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.

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

CR Number	Description
Summit X440 Series Switches	
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.
Summit X460-G2 Series Switches	
xos0061486	Combo ports have unsupported autonegotiation and half-duplex settings.
xos0062425	On Summit X460-G2 series switches, the primary port is incorrectly set as 40 when it should be 41. Under certain conditions, this can cause a kernel crash.
Summit X670 Series Switches	
xos0061167	Links become active without a connection with tri-speed Base-T SFP installed.
Summit X670-G2 Series Switches	
xos0061661	Multicast add failures occur on Summit X670-G2 stacks with the following error message: <pre> Erro: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 </pre>

Resolved Issues in ExtremeXOS 15.7.2

The following issues were resolved in ExtremeXOS 15.7.2. ExtremeXOS 15.7.2 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, ExtremeXOS 15.6.2, and ExtremeXOS 15.7.1. 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
General	
xos0062018	IPv6 ping between peers fails when IPv6 global unicast addresses are used as the default gateway.
xos0062016	Command line process memory leak occurs when accessing switches with Chalet.

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

CR Number	Description
xos0061835	The process exsh causes excessive CPU utilization when it does not terminate after performing continuous Telnet/SSH.
xos0061379	SNMP get does not retrieve the correct temperature.
xos0059655	Error "Unable to delete permanent entry" appears while deleting static blackhole FDB entries.
xos0051961	Unable to block IPv6 traffic from SSH/Telnet/Web interface by access-profile policy.
xos0057211	Traffic gets forwarded for blackholed MAC address when limit learning is enabled.
xos0057462	SSL certificate generated in switch uses a weak hashing algorithm, MD5, which is known to be vulnerable to collision attacks.
xos0057463	The minimum key length supported when configuring an SSL certificate is 64 bits, which is considered medium strength and could be exploited by an attacker on the same physical network.
xos0057599	After giving gateway as the same interface address on VLAN in <code>config bootprelay ipv6 prefix-del snooping add</code> command, error does not occur and the same route is added in snooping table and rtmgr table.
xos0057672	The process rtmgr ends unexpectedly when disabling GRE tunnels.
xos0057679	Account user name and password are not encrypted in logs when cli-config-logging is enabled.
xos0058241	XNV Dynamic VLAN: After disabling, and then enabling, vm-tracking, the <code>show vm-tracking</code> command displays the incorrect number of network VMs authenticated with NMS as the authentication method.
xos0058349	Enabling DHCP on any VLAN with IP address already configured does not produce error.
xos0058668	After rebooting DHCPv6, client remains in rebooting state.
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> .
xos0059952	XNV Dynamic VLAN: After disabling dynamic VLAN on vmt-enabled port, the "v" flag is not removed from the output of the <code>show fdb</code> command.
xos0059989	Configuring non-persistent command using UPM script shows dirty bit(*) in the prompt.
xos0060088	Kernel oops triggered rarely during continuous addition/deletion of ARP entries for long duration in presence of high CPU utilization.
xos0060092	Fetching values using SNMP for "extremePortQosStatsTable" does not work correctly.

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

CR Number	Description
xos0060119	Changing the primary TACACS server configuration locks out TACACS-authenticated users.
xos0060228	HAL process ends unexpectedly in rare circumstances while rebooting switch with default speed configuration on 10G ports.
xos0060354	ExtremeXOS ScreenPlay using IPv6 does not work with HTTPS.
xos0060713	Executing the command <code>clear cpu-monitoring</code> sets the dirty bit(*).
xos0060825	Double-tagged CFM frames are dropped by kernel in VMAN environment.
xos0060909	In UPM profiles the variable <code>EVENT.TIME</code> incorrectly has the current time rather than the time when the event was queued/triggered.
xos0060977	MPLS traffic is egressing switch with incorrect label causing packet loss.
xos0061085	Kernel oops occurs while deleting VR with enable BGP export and IPARP proxy configurations.
xos0061517	LACP adjacency fails while forwarding the PDU with l2pt profile over L2VPN tunnels when MPLS PHP is enabled.
xos0061565	The TCL function, "clock scan," generates errors with default time zone configuration.
xos0061656	Nodes remain in the "FDBSync" state due to temp-flooding while rebooting the stack.
xos0061699	Traffic is dropped when moving idmgr client from one port to another with role-based authentication.
BlackDiamond 8800 Series Switches	
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.
xos0061822	HAL process ends unexpectedly during failover when switches have ACL policies without meter action.
xos0061338	Packets are not switched on port-specific, tag-enabled ports on XL series I/O modules.
xos0060891	While sending continuous join and leave, as well as multicast streams multiple times allowing the streams to age out, kernel error messages appear for add/delete operation with reason "Entry not found".
xos0058950	XFP optics appear as unsupported in the output of the <code>show port configuration</code> command.
xos0059648	Static ARP entries are not properly synced with new Master Switch Fabric Module after failover.
xos0059789	Dos-Protect ACL is not cleared after continuous DOS attack occurs.

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

CR Number	Description
BlackDiamond X8 Series Switches	
xos0057827	Layer 2 multicast traffic is not forwarded to IGMP receivers after MM (management module) failover.
xos0058842	Tagged VLAN traffic is dropped at ingress of 100G4X-XL module port, when the same port is added to another CEP-VMAN VPLS. This issue does not occur on non-100G4X-XL modules/ports.
xos0060210	Rarely, HAL process may end unexpectedly due to buffer overflow condition while running diagnostics for Management Module.
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.
xos0061902	BlackDiamond X8 series switches use VLAN instance as index instead of router interface (rtif) for ARP entries.
Summit Family Switches	
xos0059248	Identity management: Identity management ends unexpectedly with signal 11 repeatedly when NetLogin MAC addresses are scaled to 512-3,000.
xos0059950	In Summit series switches, you cannot download bootROM images from memory card.
xos0060258	After failover, ports are not added to aggregator in LACP with passive mode. Port are added correctly when switches are in active mode (default state).
SummitStack	
xos0061957	HAL process ends unexpectedly during failover when switches have ACL policy without meter action.
xos0057438	Memory depletion occurs in Backup/Standby nodes of SummitStack with highly scaled IPFIX flow records.
xos0057835	In SummitStacks, clear-flow sampling period is incorrectly calculated.
xos0059462	Timezone configuration is not applied to standby nodes after stack reboot.
xos0059434	Unable to re-initialize MAC-authenticated client using extremeMacAuthClientInitialize OID.
xos0060115	Expected <Info:Kern.Info> log messages are not shown for log filter events.
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.
xos0061479	Openflow process signal 6 ends unexpectedly on Summit X460 stacks.

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

CR Number	Description
Summit X430 Series Switches	
xos0059934	Downloading BootROM corrupts the bootloader.
xos0057028	In Summit X430 series switches, kernel gets stuck for few seconds while installing boot v1.0.1.5 from ExtremeXOS.
Summit X440 Series Switches	
xos0050402	The command <code>enable inline-power legacy</code> does not power up pre-standard PoE devices, such as Cisco phone 7940/7960 that do not work with IEEE 802.3af standard detection and legacy capacitive detection. The <code>enable inline-power legacy</code> command now powers up legacy PoE devices that rely on the capacitive detection instead.
xos0054084	In Summit X440 stacks with about 500 identities, restarting the <code>idmgr</code> process, causes the <code>IdMgr</code> process to end unexpectedly with signal 11.
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.
Summit X460 Series Switches	
xos0058250	Installing an ExtremeXOS image on a SummitStack after copying the image to the switch using SCP fails when there is no backup module assigned to the stack.
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 stacks with mixed alternate and native stacking enabled slots, traffic ingressing one specific slot is not forwarded to other slots.
Summit X460-G2 Series Switches	
xos0060998	MAC-based authenticated sessions over RADIUS are erroneously deleted by FDB.
xos0060736	gPTP propagation delay is not calculated correctly and ports become AVB incapable.
xos0061283	Links do not come up on the last four ports (49-52).
xos0061313	Summit X450/X460-G2 series switches 10G ports flood unicast traffic to other ports on the switch.
Summit X670 Series Switches	
xos0061770	Detected parity errors may cause a kernel crash.
xos0055971	In Summit X670v-48x switches, the dirty bit (*) is set if there are 1000BaseT SFPs installed.

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

CR Number	Description
xos0061730	Traffic is affected when unconfiguring ACL on another port.
Summit X670-G2 Series Switches	
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.
Summit X770 Series Switches	
xos0055746	Stacking port link flap occurs on Summit X770 series switches when using 3-meter QSFP+ cables.
E4G-200 Cell Site Routers	
xos0060715	SNMP get does not return the product name for E4G-200-12x cell site routers.
ACL	
xos0054348	Cannot delete flow names after deleting, and then creating, the flow while the ACL is installed.
xos0061922	Dynamic ACLs applied as "any" fail to install in hardware after upgrading ExtremeXOS from any release other than EXOS 15.3.
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.
BGP	
xos0060407	After disabling, and then enabling EBGP, new routes from different autonomous systems (ASs) are not considered for best path calculation..
xos0061310	DCBGP process ends unexpectedly after disabling BGP peer switches after the switch receives VPNv4 routes with the same next hop from two different L3VPN peers.
EAPS	
xos0061038	Loops occur in EAPS-protected VLANs, after peer reboot, if a VLAN's port is also protected by ELSM.
ERPS	
xos0059320	CCM is dropped for "Hardware Down MEPs" when they are received on ports that are blocked by ERPS.
ESRP	
xos0061965	Configuring ESRP member VLANs (VRRP-enabled) produces errors.

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

CR Number	Description
MLAG	
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.
xos0061586	In two-tier MLAG topology, loops occur if the VLAN header has the CFI bit set.
NetLogin	
xos0061069	In NetLogin ISP mode, client MAC addresses configured as static FDBs are removed after reboot.
xos0060633	Possible memory leak while cleaning the authVLANlist in NetLogin during unauthorization.
xos0060449	NetLogin MAC-based mode does not work as expected with PVLANS.
xos0060242	Ports are not removed when unauthenticated UPM profiles are triggered.
xos0053634	MAC-lockdown-timeout on user ports does not work as expected if NetLogin is enabled on those ports.
xos0056455	Creating NetLogin local user with encrypted option does not create the user because the password is not copied if encrypted option is specified.
OSPF	
xos0056243	OSPF process ends unexpectedly during frequent route recalculation caused by switch reboot or MSM failover or BFD flap events.
xos0058221	Rarely, OSPFv3 process ends unexpectedly with signal 11 when link flaps occur.
xos0059574	OSPF packets larger than 8,192 are dropped even with jumbo frame enabled.
Policy	
xos0061178	Dynamic ACL for gratuitous ARP violation on LAG member ports are incorrectly getting installed on LAG master ports.
Python	
xos0060600	EPM process signal 5 ends unexpectedly when creating a new process that has dash.
SNMP	
xos0059964	SNMP poll for MIB dot3StatsDuplexStatus always returns unknown(1) when ports are configured with auto-negotiation on.

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

CR Number	Description
TWAMP	
xos0060780	With VRRP enabled, local VLAN's direct route is not installed in hardware after reconfiguring the VLAN's IP address.
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.
VLAN	
xos0060103	SNMP walk does not return all VLANs under extremePortVlanStatsTable.
xos0054039	IP multicast traffic is not forwarded on PSTAG VLANs when it shares ports with other IGMP snooping-enabled VLANs or other L3 VLANs.
VRRP	
xos0061222	Gratuitous ARP packets for VRRP virtual IP addresses have ARP sender addresses as physical MAC addresses, instead of VRRP virtual MAC addresses.
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.
xos0053821	IPv6 neighbor advertisements for VRRP virtual IP address uses virtual MAC address as source MAC address instead of switch MAC address.

Resolved Issues in ExtremeXOS 15.7.1

The following issues were resolved in ExtremeXOS 15.7.1. ExtremeXOS 15.7.1 includes all fixes up to and including ExtremeXOS 11.6.5.3, and earlier, ExtremeXOS 12.0.5, ExtremeXOS 12.1.7, ExtremeXOS 12.2.2-patch1-12, ExtremeXOS 12.3.6, ExtremeXOS 12.4.5, ExtremeXOS 12.5.5, ExtremeXOS 12.6.3, ExtremeXOS 12.6.5, ExtremeXOS 12.7.1, ExtremeXOS 15.1.5, ExtremeXOS 15.2.4, ExtremeXOS 15.3.3, ExtremeXOS 15.4.1, ExtremeXOS 15.5.1, ExtremeXOS 15.5.2, ExtremeXOS 15.6.1, and ExtremeXOS 15.6.2. 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
General	
xos0058661	ARP reply packets are dropped with continuous IP unicast miss packets with DSCP 48.
xos0057421	Fixed vulnerability in OpenSSL 1.0.1 through 1.0. that can affect: <ul style="list-style-type: none"> • THTTPD • Tech Support feature • Optics licensing feature • OpenFlow • XMLC • CLI master one-time debug password feature
xos0056340	Unknown Layer 2 traffic from Isolated subscriber VLANs are forwarded to the remote MLAG ports, even though local MLAG ports are up.
xos0056987	With ELSM enabled on port p1 on two switches, with 10 seconds as the ELSM hello interval, and trap receiver configured on the both switches (or managed in Ridgeline, which configures the trap receiver by itself), and then disable/enable port p1 on both switches. Switches send two linkup traps.
xos0054199	Ingress traffic stalls on port when switches receive continuous 802.3x pause frames on egress ports for that traffic stream.
xos0055108	The bound IP address is not being reflected in the command <code>show vlan</code> .
xos0059851	When a DHCP client receives an IP address that conflicts with a static IP address configured in a VLAN, the static IP address is removed from the VLAN and the DHCP client stops.
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.

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

CR Number	Description
xos0056683	CCM sender TLV not recognized due to switches using different SenderID TLVs to convey the same information (IP address).
xos0057407	Hops fields in DHCP packets are not incremented when processed by Bootprelay.
xos0057601	The command <code>disable ip-security arp learning learn-from-arp vlan <vlan_name> port <portname></code> is not saved/applied after reboot.
xos0055398	Authentication fails for a Netlogon client in dot1x mode, since the port added untagged in one VLAN cannot be moved to another VLAN.
xos0057199	The output of the command <code>ls <filename></code> displays the last accessed time stamp instead of the last modified time stamp.
xos0053350	The TFTP put and get commands do not support the current working directory operaoor (.) in their arguments. The default working directory is always taken as the <code>/cfg</code> directory irrespective of the current directory from which the command is issued.
xos0057606	"enable cli prompting" is allowed inside UPM/CLI scripts.
xos0058393	The TCL command <code>clock format</code> is not available for CLI scripting and attempting to use it produces an error.
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 TAB .
xos0056191	Executing the command <code>show temperature</code> from a user-created account returns an error.
xos0056228	TFTP get operation fails when the remote file exists in second level sub-directory.
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.
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.
xos0058968	Error log "Function Pointer Database is not fully initialized" appears during bootup on non-Summit platforms.
xos0053046	The command <code>delete meter <meter-name></code> produces the error "Error: Timeout awaiting meter deletion" and the meter is not removed from master switch, but the meter is removed from backup switch.

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

CR Number	Description
BlackDiamond 8800 Series Switches	
xos0054970	BlackDiamond 8800-xl cards and Summit X480 series switches should not allow Layer 2 Protocol Tunneling and Filtering to be configured over VPLS/VPWS.
xos0059605	Sys-health-check output shows false fabric port flap events between Master Switch Fabric Module (MSM) and I/O module.
xos0057624	On a PVLAN, after restarting process VRRP on the VRRP master switch, traffic loss occurs in the switch.
BlackDiamond X8 Series Switches	
xos0053636	Changing the IP MTU on VLANs, causes rtmgr process to end unexpectedly.
xos0054860	Configuring, and then unconfiguring, large ACL policy files two or three few times causes the ACL to fail during subsequent installs.
xos0054861	Incorrect frame delay measurements occur on BlackDiamond X8 and 8800 series switches since kernel time stamping for software end points isn't supported.
xos0055433	The tDiag process occasionally ends unexpectedly after executing <code>show debug system-dump MM B from MM-A</code> , when MM-B does not contain system dump.
xos0056184	Removing primary Management module (MM) causes LAGs to bounce.
xos0056300	Traffic does not switch back to primary port when smart redundancy is enabled.
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><Error:DM.Error> MM-A: Process thttpd Failed MM-A rebooted</code>
xos0057643	On Summit X460-G2, X770, and BlackDiamond switches, slow learning rates occur for scaled multicast traffic in I2-and-I3-and-ipmc mode.
xos0058375	ACLs to match VLAN-ID, CVID parameters do not work for slow path forwarded packets.
xos0058568	Some front panel ports cannot be enabled after rebooting the I/O module.
xos0058847	After second Master Switch Fabric Module (MSM) failover, backup root gets has incorrect STP bridge ID.
xos0059104	ACL policies are not installed in hardware after management module failover.
xos0059156	VRRP control packets are dropped due to congestion in tx queue under scaled environments.

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

CR Number	Description
xos0059343	The process snmpMaster might end unexpectedly during upgrade from ExtremeXOS 15.3 to 15.5 for some SNMP community names.
xos0059733	LSP load sharing does not occur on Summit X460-G2 and BlackDiamond X8-100G4X switches.
Summit Family Switches	
xos0052494	RSTP port transition takes longer than usual to go to forwarding state if STPD is disabled, and then enabled.
xos0053409	After enabling third-party optic license, traffic should be forwarded in 40G and log messages about third-party optics should not occur.
xos0053755	Should have same session ID appear in both <code>show session</code> command and log target field in <code>show log configuration</code> command.
xos0055637	With two flows installed on the a Summit X440 series switch, only one flow is recognized.
xos0056230	SNMP query on "extremeMemoryMonitorsystemTable" does not show backup information, if slot2 is master and slot1 is backup.
xos0056296	Make p2p calculations more robust against negative path delay measurements.
xos0056704	Cannot overwrite dynamic unicast addresses with static multiport unicast addresses.
xos0056812	Non-default configuration for mac-lockdown-timeout is removed from the configuration when mac-lockdown-timeout is disabled on a port.
xos0057106	When mirroring is configured to be triggered through clearflow, mirroring does not work and produces the following error: 06/26/2014 17:14:16.40 <Warn:HAL.IPv4ACL.Warning> : Could not enable mirroring for ACL rule since mirror acl-rule-1 is not active.
xos0057897	Orphan static route not preserved after reboot when ECMP route exists.
xos0058040	DNS name resolve fails when name-server is first added to VR-default, then to VR-mgmt, and then deleted from VR-default.
xos0058048	AVB data traffic sent with dot1p "2" priority at wire-rate causes gPTP port state to flap.
xos0058050	After 20 minutes of running audio, noise periodically occurs.
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.

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

CR Number	Description
xos0056878	Summit X440 series switches are not inserting the MSTI Configuration Message (MSTI Vector information for MSTI instance 1) into BPDU packets when configured as the root bridge for both the CIST (stpd s0) and for MSTI 1 (stpd s1) in region 1.
SummitStack	
xos0053998	In Summit X670 stacks, the following error messages appear after executing the command <code>debug hal show optic-info slot 3</code> , and then saving and rebooting: <Error:HAL.Port.Error> Slot-3: getQSFPInfo(748) Error reading eeprom <Warn:Kern.Card.Warning> Slot-3: i2c-1: shid_eeprom_read:3244 I/O failed, addr 0x0050 rd/wr 0 cmd 127 proto 1, rc 145 This is not observed with 15_4_0_40 build
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.
xos0057767	Static FDB associated with VPLS service VLAN is not programmed in hardware after reboot when <code>disable learning</code> is configured.
xos0058142	System crashes with “process nodemgr pid 1595 signal 5” error when master slot is rebooted with maximum number of port-based mirroring filters configured and multicast traffic.
xos0058430	SummitStack firmware versions reported by SNMP are not valid.
xos0056179	Commands use inconsistent syntax for value ranges, some of which are incorrect and/or misleading.
xos0058133	SummitStacks use slot MAC address in DHCP packets when enabling DHCP on Management VLAN.
Summit X430 Series Switches	
xos0057663	For Summit X430 series switches, the command <code>show inline-power</code> shows incorrect value for budgeted power as set in the command <code>config inline-power budget</code> .
xos0059524	Link status is incorrect when auto-polarity setting is off.
Summit X440 Series Switches	
xos0058068	Summit X440-24tDC switches are reporting maximum temperature limit 60°C under normal conditions.
xos0058184	When the designated router (DR) receives a register-stop message from the rendezvous point (RP), it starts a register-stop timer to maintain this state. Just before the register-stop timer expires, the DR sends a null-register message to the RP to allow the RP to refresh the register-stop information at the DR. Summit X440 series switches acting as DRs do not send null register before the register-stop timer expires.

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

CR Number	Description
xos0058300	Packets are dropped on combo ports when the preferred medium is configured as copper force.
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").
xos0056738	Summit X440 series switches do not drop the IGMP Query Message with an invalid IGMP checksum.
Summit X450 Series Switches	
xos0057647	Packets are forwarded to CPU after deleting the VLAN with disable learning.
Summit X460 Series Switches	
xos0051266	In Summit X460 stacks, backup does not get configuration synchronized, when forming a stack with only one stacking port using alternate stacking mode.
xos0055518	On Summit X460 series switches with OpenFlow enabled, switches boot up with the following error messages: <pre> "Error while loading structure <cfgTechSupport><basic>1</basic><hour>0</ hour><srcipaddress>10.68.63.86</ srcipaddress><hostname><![CDATA[12.38.14.200]]></ hostname><sslEnabled>0</sslEnabled><port>800</ port><daily>0</daily><isDefault>1</ isDefault><bootup>1</bootup><automatic>1</ automatic><nameIndex><![CDATA[12.38.14.200]]></ nameIndex><errorDetected>0</ errorDetected><vrName><![CDATA[VR-Mgmt]]></ vrName></cfgTechSupport>: Source IP address 10.68.63.86 does not belong to the VR VR-Mgmt. <cfgTechSupport>" "Error while loading structure <openflowGlobalConfig><numVlans>2</ numVlans><version>0</version><fdb>0</ fdb><isEnabled>1</isEnabled></ openflowGlobalConfig>: Error: Invalid OpenFlow version<openflowGlobalConfig>" </pre>
xos0056342	Misleading power supply unit (PSU) traps are sent when PSUs are inserted or powered on/off.
xos0058043	MSRP: packets are dropped when bandwidth is increased to 611 M.
xos0058053	Summit Stack run failover takes longer than usual time to boot the backup node.

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

CR Number	Description
xos0058248	Summit X460 series switches change boot image after Installing an xmod image.
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.
xos0059671	On Summit X460 series switches with 750 W power supplies installed, log messages "Power usage data unknown" appear.
Summit X460-G2 Series Switches	
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.
xos0057245	In Summit X460G2-24p-G4, 24t/p, and 48t/p (1G types), links do not come up for 100 Mbps and 10/100/1000 Base-T optics.
xos0059240	High CPU utilization and MAC learn thrashing occurring on fabric links forming a X460-G2 stack.
xos0059577	On Summit X460G2 series switches, can't install ExtremeXOS SSH XMOD image.
xos0057346	Link flaps occur when optics such as 10/100/1000 Base-T or 100FX is inserted into Summit X460-G2 and X480 series switches.
xos0056971	In Summit 670G2-72x and Summit 460-G2 stacks, executing a failover produces the following Kernel error: <code>Errno:Kern.Card.Error> Slot-1: KICM: Could not change SP1 type for NONE neighbor. idx1=0,idx2=0</code>
Summit X670 Series Switches	
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.
xos0059128	In Summit X670 series switch, all LEDs are blinking at a faster rate.
Summit X770 Series Switches	
xos0053986	Packet buffer cells are not drained out when fabric flow control is disabled in a stack and also when PFC RX pause is disabled/enabled during user stream and PFC frames are active in the switch.
xos0059573	Factory installed image incorrectly references "x450" in the image name.

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

CR Number	Description
AVB	
xos0058603	PTP follow up does not happen correctly when correction field is greater than 32 bits.
ACL	
xos0054720	With network-zone configured, the command <code>show access-list port 1:1 detail</code> reverses the IP address match criteria.
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.
xos0059330	With dual master switch fabric module (MSM) installed, clear-flow ACL intermittently fails.
ERPS	
xos0056122	Configuring MEP IDs through ERPS (for dynamic CFM creation) commands are unavailable.
xos0057141	Default VLAN cannot be added as a protected VLAN to an ERPS ring.
xos0058464	In ERPS rings, blocking the control channel by deleting the ports from the control VLAN causes a short loop in the ring.
xos0057179	ESRP feature is not enabled immediately after the installing an Advance Edge license.
FDB	
xos0059146	With port-specific tags configured, source MAC addresses are removed and re-learned for all incoming ARP packets causing flooded traffic a for short time interval.
IP General	
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.
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.
xos0059581	Rtmgr process ends unexpectedly when OSPF external routes are deleted from the route table.
IP MLD	
xos0054105	Source IP addresses for MLDv1 hosts in user-defined PIM SSM range are not resolved.
xos0054121	Problem with storing SSM mapping entries with SSM mapping entries in hierarchal manner.

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

CR Number	Description
xos0057688	The command <code>clear ipmc fdb</code> does not clear IPv6 multicast entries.
IP PIM	
xos0053818	System stops responding when PIM is disabled for the ingress interface an on MLAG peer when there is a forwarding egress interface and at least another one asserted.
xos0054101	W-MLAG-PIM: (*;G)s should not be created on non-designated router (DR) after PIM is bounced. Ingress interface for PIM DR is an MLAG VLAN with clients present.
MLAG	
xos0053412	After peer reboots, MLAG ports remain in ready state.
xos0053469	MLAG health check: Hello packets are exchanged temporarily when the alternate path is restored after ISC connectivity is lost.
xos0053744	MLAG health check: Manually disabled MLAG ports that are part of static/LACP sharing are re-enabled by alternate health check.
xos0058873	FDB entries are learned incorrectly on VMANs in MLAG peers when the MLAG ports are CEP ports for multiple VMANs with different CVIDs.
MPLS	
xos0058880	Packets are not switched to primary path after recovering from path failure in MPLS RSVP-TE.
xos0056994	Unable to add EAPS shared ports to VLANs even after disassociating them from VPLS domains.
OpenFlow	
xos0058117	After enabling, and then disabling, OpenFlow globally and on a VLAN with standard mode, you cannot enable learning on the VLAN anymore.
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.
xos0055946	In 5-node setup, after installing the OpenFlow groups and flows using OVS commands, the transit node is not forwarding traffic to all applicable switches.
xos0056061	The OpenFlow switch is not actively sending out hello messages when the 3-way TCP connection is established.
OSPF	
xos0058056	OSPF-opaque, LSA-related configurations do not appear in output of the <code>show configuration</code> command.
xos0059305	OSPF consumes a large amount of memory when a large number of Link State Acknowledgment packets are queued up for transmission.

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

CR Number	Description
PoE	
xos0058473	Operator-set power is rounded down to the nearest 1.0 Watt. Ports display the configured maximum allocated power, rather than the actual power. This can lead to power being denied to devices requesting power under the configured limit.
xos0058994	POE is not delivering power to several model phones when legacy mode is enabled.
Python	
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 <code>./config/xos_config.xml</code> : no such file or directory."
xos0058835	Python scripts can disable command line output after the script finishes.
RIP	
xos0058683	RIP packets are dropped when another VLAN has a secondary IP address configured.
STP	
xos0059002	Checkpoint errors occur during execution of STP debug command if switch contains many STP-enabled VLANs.
VLAN	
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.
xos0057435	Packets are dropped when learning is disabled in a VLAN when its associated ports are configured with limit learning in another VLAN.

4 ExtremeXOS Documentation Corrections

This chapter lists corrections to the *ExtremeXOS 15.7 User Guide* and *ExtremeXOS 15.7 Reference Guide*.

This chapter contains the following sections:

- [ACLs on page 141](#)
- [ACL Egress Counters Limitation on page 141](#)
- [ACL Policy Redirect on page 142](#)
- [ACL Ports Limits on page 142](#)
- [BOOTPrelay Not Supported in Virtual Routing and Forwarding \(VRF\) on page 143](#)
- [Configure Access-List VLAN-ACL-Precedence Command Usage Guidelines on page 144](#)
- [Configure IP-MTU VLAN Command Syntax Description on page 145](#)
- [Configure Sys-Recovery-level Slot Command Platform Availability on page 146](#)
- [Debounce Commands on page 147](#)
- [ELRP on page 149](#)
- [ELRP and QoS on page 149](#)
- [Extreme Networks Virtualization \(XNV\), Identity Management \(IDM\), and Network Time Protocol \(NTP\) on page 150](#)
- [Hardware Table Hash Algorithm on page 152](#)
- [L2VPN Sharing Commands on page 151](#)
- [Link Aggregation \(LAG\) Limit for Multiprotocol Label Switching \(MPLS\) Terminated Packets on page 152](#)
- [Link Layer Discovery Protocol \(LLDP\) on page 153](#)
- [MIB on page 153](#)
- [MLAG on page 154](#)
- [MLAG PIM-SM *,G Forwarding Limitation on page 154](#)
- [NetLogin Limitation on page 155](#)
- [NetLogin Local Authentication on page 155](#)
- [PoE Power Delivery on page 156](#)
- [PVLAN + EAPS and PVLAN + STP Recommendations on page 157](#)
- [Rate Limiting/Meters on page 158](#)
- [Remote Mirroring on page 159](#)
- [Routing Policies on page 160](#)
- [Synchronize Command on page 161](#)
- [TACACS Server on page 161](#)

- [VRRP Guidelines on page 163](#)

ACLs

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

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

xos0062545

The following content:

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

Needs to be changed to:

“redirect-port-list : port_list—Supports multiple redirect ports as arguments. When used in an ACL, matching packets are now redirected to multiple ports as specified in the ACL while overriding the default forwarding decision. Maximum number of ports that can be mentioned in this list is 64. (Summit X440, X460, X480, X670, X770, E4G-200, E4G-400, BlackDiamond 8K - 8900-G96T-c, 8900-10G24Xc, 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 Guide for the `configure sys-recovery-level slot` command.

xos0061985

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

“Platform Availability

This command is available only on modular switches.”

Should change to:

“Platform Availability

This command is available on modular and stacking switches.”

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

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:

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.

Extreme Networks Virtualization (XNV), Identity Management (IDM), and Network Time Protocol (NTP)

ExtremeXOS User Guide, ExtremeXOS Feature License Requirements

xos0057855

The L2 Edge license no longer supports Extreme Networks Virtualization (XNV), Identity Management (IDM), and Network Time Protocol (NTP) features from ExtremeXOS 15.5.2 onwards for the Summit X430 series switches.

- 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>
```

L2VPN Sharing Commands

ExtremeXOS Command Reference Guide, L2VPN commands

xos0062252

Remove the following commands:

- *configure l2vpn sharing hash-algorithm*
- *configure l2vpn sharing ipv4*

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

- *enable l2vpn sharing*
- *disable l2vpn sharing*

Platform Availability:

This command is available only on the following switches

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

Hardware Table Hash Algorithm

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

xos0061737

The following text:

“Use the following command:

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

Should change to:

“Use the following command:

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

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 *LLDP chapter > Configuring and Managing LLDP > Enable and Disable LLDP*

xos0061698

The existing text:

“LLDP is disabled on all ports by default.”

Should change to:

“LLDP is enabled on all ports by default.”

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

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

The existing text:

“Default: Disabled”

Should change to:

“Default: Enabled”

MIB

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

xos0062086

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

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

MLAG

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

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

xos0062552

Add the following note:



NOTE

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

NetLogin Limitation

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

xos0062943

Add the following additional limitation:

“When using Netlogin MAC-based VLAN mode, moving a port as untagged from the 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”

PoE Power Delivery

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

xos0062239

Add the following note:

**NOTE**

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

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

PVLAN + EAPS and PVLAN + STP Recommendations

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+.

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 is 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

